Local Observation and Global Objecthood: Sheaf Descent, Gluing Obstructions, and the Algebra of Distributed Concept Formation

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

July 10, 2026

Mathematical objects are often constructed locally before they are recognized globally. A sensor observes one region, a classifier resolves one subset of data, a scientific discipline stabilizes identities within a limited context, and a mathematical chart describes only part of a larger space. The central problem is therefore not merely how an observation policy forms quotient objects, but when many local policies can be combined into one coherent global system of identity. This paper develops a local-to-global theory of object formation using restrictions, overlaps, generated equivalence relations, descent conditions, and gluing obstructions. Compatible local equivalences do not always determine a faithful global equivalence, because chains of identifications can create new relations that no local observer explicitly licensed. The resulting transitive closure may produce emergent global objects, destructive fusions, or false gluing. We define minimal globalization, faithful descent, gluing defects, and transformation-stable global quotients, then organize these constructions categorically. The framework extends the algebra of conceptual resolution from single-policy state spaces to distributed observation networks and provides foundations for sensor fusion, artificial intelligence, scientific integration, multiscale modeling, biological classification, and contextual physical theories.

Keywords: local-to-global object formation, distributed observation, conceptual resolution, observation policies, equivalence relations, quotient objects, sheaf theory, descent theory, gluing obstruction, false gluing, faithful descent, equivalence closure, contextual identity, observer networks, stable quotients, category theory, artificial intelligence, sensor fusion, multiscale modeling, scientific integration.

A collaboration with OpenAI GPT-5.5 Thinking. CC4.0. 239 pages and 62 references.

https://orcid.org/0009-0004-4206-4253

Outline

1. Introduction: From Individual Observation to Distributed Objecthood

1.1 The stable-quotient foundation

1.2 Why a single observation policy is insufficient

1.3 Local knowledge and partial object formation

1.4 The local-to-global problem

1.5 Agreement on overlaps versus global coherence

1.6 Transitivity as the hidden gluing problem

1.7 Principal contributions and scope

1.8 Main results

2. Local State Spaces and Restricted Observation Policies

2.1 Global state spaces and observational domains

2.2 Covers by local regions

2.3 Observation policies on subsets

2.4 Restriction of equivalence relations

2.5 Local quotient object spaces

2.6 Refinement within observational domains

2.7 Restriction maps between resolution lattices

2.8 Examples of locally formed objects

3. Overlap Compatibility and Descent Data

3.1 Pairwise intersections of observational domains

3.2 Agreement of identity policies on overlaps

3.3 Compatible families of local equivalences

3.4 Cocycle-like consistency on multiple overlaps

3.5 Local object correspondences

3.6 Why pairwise agreement appears sufficient

3.7 The difference between compatibility and descent

3.8 Descent data for distributed observation

4. Generated Globalization and Equivalence Closure

4.1 Union of local identity relations

4.2 Failure of the union to be an equivalence relation

4.3 Reflexive, symmetric, and transitive closure

4.4 Minimal equivalence generated by local policies

4.5 Zig-zag paths of identification

4.6 The minimal globalization theorem

4.7 Universal property of generated globalization

4.8 Global quotient objects produced by closure

5. The Faithful Descent Theorem

5.1 Exact recovery of local policies

5.2 Restricting the generated global relation

5.3 Statement of the faithful descent criterion

5.4 Necessity of local recovery

5.5 Sufficiency of the restriction condition

5.6 Proof of the faithful descent theorem

5.7 Uniqueness and minimality of faithful globalization

5.8 Interpretation as lawful global object formation

6. Gluing Obstructions and Descent Defects

6.1 Identification chains crossing observational domains

6.2 New local relations generated globally

6.3 The gluing obstruction relation

6.4 Quantifying descent defects in finite systems

6.5 Obstruction witnesses and minimal zig-zags

6.6 Local consistency without global fidelity

6.7 Repair by refinement, separation, or cover modification

6.8 An obstruction theory for conceptual integration

7. False Gluing, Failed Gluing, and Emergent Identity

7.1 Three distinct outcomes of globalization

7.2 Faithful gluing

7.3 Failed gluing through incompatibility

7.4 False gluing through excessive closure

7.5 Emergent global identity

7.6 Destructive fusion of locally distinct objects

7.7 Distinguishing emergence from observational error

7.8 The ontology of globally generated objects

8. Stable Dynamics on Glued Quotient Spaces

8.1 Local admissible transformation monoids

8.2 Compatibility of transformations on overlaps

8.3 Global transformations assembled from local actions

8.4 Invariance of generated equivalence

8.5 The stable globalization theorem

8.6 Induced dynamics on the global quotient

8.7 Local stability versus global stability

8.8 Dynamical failure of distributed abstractions

9. The Presheaf of Conceptual Resolutions

9.1 Observation policies assigned to domains

9.2 Contravariant restriction along inclusions

9.3 The resolution presheaf

9.4 Identity and composition of restriction

9.5 Failure of the ordinary sheaf condition

9.6 Separatedness and uniqueness of globalization

9.7 Sheafification as enforced conceptual closure

9.8 What sheafification preserves and destroys

10. Categorical Descent and Observer Networks

10.1 Categories of local policy-indexed spaces

10.2 Restriction functors

10.3 Diagrams indexed by observational covers

10.4 Cones, limits, and compatible observer systems

10.5 Descent objects and global realizations

10.6 Morphisms between distributed observation networks

10.7 Refinement and coarsening across networks

10.8 Toward stacks and higher observational coherence

11. Examples and Computational Constructions

11.1 A three-domain finite gluing example

11.2 A compatible family that fails faithful descent

11.3 False fusion produced by a transitive zig-zag

11.4 Distributed image classification

11.5 Multisensor object tracking

11.6 Biological taxonomies assembled across laboratories

11.7 Multiscale physical descriptions

11.8 Algorithms for detecting gluing defects

12. Distributed Object Formation Theorem and Research Program

12.1 Statement of the combined local-to-global theorem

12.2 Compatible local policies and minimal globalization

12.3 Faithful descent and its obstruction

12.4 Stable transformation on global quotients

12.5 Presheaf structure and sheaf failure

12.6 Implications for artificial intelligence and science

12.7 Limits of set-theoretic local identity

12.8 Next steps: graded descent, stacks, and computational discovery

13. References

1. Introduction: From Individual Observation to Distributed Objecthood

Mathematical object formation does not always occur within a single observational domain. Many objects are assembled from partial views, local measurements, restricted classifications, or context-dependent descriptions. A telescope records one region of the sky, a microscope resolves one field of tissue, a laboratory studies one population, a sensor observes one sector of an environment, and a mathematical chart describes one portion of a space. No individual observation necessarily contains the whole object. Global identity must instead be reconstructed from a collection of local identities.

This creates a problem that does not arise when one observation policy acts on one fully accessible state space. In the single-policy setting, an equivalence relation determines which states count as the same, and the corresponding quotient produces the effective objects visible at that resolution. Stability under admissible transformation ensures that the quotient supports representative-independent reasoning. The central issue is whether the chosen equivalence relation is preserved.

Distributed observation introduces an earlier difficulty. Before asking whether a global quotient is stable, one must determine whether a coherent global quotient exists at all.

Suppose a state space is covered by overlapping observational domains. Each domain carries its own equivalence relation. Within each region, the local observer can decide which states count as the same. On overlaps, observers can compare their decisions. If their equivalence relations agree wherever their domains overlap, it may appear that the local policies can simply be combined into a global policy.

That conclusion is not generally valid.

Equivalence relations are transitive. Local identifications can combine across a chain of overlapping domains and generate additional identifications that were not explicitly licensed by any one observer. A state in one region may be identified with a state in a second region, which is identified with a state in a third region, thereby forcing a global identification between states that were never jointly observed. When the resulting global relation is restricted back to a local domain, it may identify states that the original local policy distinguished.

The local policies may therefore agree wherever direct comparison is possible while still failing to form a faithful global object system.

This phenomenon is the central subject of the present paper. The local-to-global problem of object formation asks when partial identity policies can be assembled without contradiction, uncontrolled fusion, or loss of locally preserved distinctions. It also asks what should be done when faithful assembly is impossible.

The simplest globalization procedure is to take the union of all local equivalence relations and then form the smallest global equivalence relation containing that union. This generated equivalence always exists. It supplies a minimal global relation compatible with all declared local identifications. Yet existence alone does not guarantee fidelity. The transitive closure required to make the relation global may introduce new identities. Globalization can therefore alter the very local policies from which it was constructed.

This distinction separates minimal globalization from faithful descent.

Minimal globalization asks for the smallest global equivalence relation containing every local equivalence.

Faithful descent asks whether that global equivalence restricts to exactly the original relation on every local domain.

The first construction is always possible. The second requires an additional condition.

The difference has broad implications. Scientific data are routinely integrated across instruments, laboratories, scales, and disciplines. Artificial intelligence systems combine classifications produced by specialized models. Sensor networks infer global objects from partial detections. Biological taxonomies are assembled from observations made in different regions and under different measurement regimes. Physical theories relate local descriptions across coordinate patches, scales, and experimental contexts. Databases merge records created under different naming and identity conventions.

In each case, global objecthood depends on gluing local judgments of sameness.

A successful gluing process creates a global identity system that preserves all locally intended distinctions. A failed gluing process exposes direct incompatibility among local policies. A false gluing process is subtler: the local policies appear compatible, but their global closure produces additional identifications that erase distinctions no observer agreed to erase.

There is also a fourth possibility. The new global identification may not be an error. It may represent an emergent object visible only through the network of local relations. A global identity can be generated by distributed evidence even though no single observer has access to the entire identification path.

The mathematics must therefore distinguish destructive fusion from legitimate emergence. Both arise through transitive globalization. The difference cannot be determined from closure alone. It depends on the admissible transformations, explanatory purposes, and identity policies governing the system.

The theory developed here extends stable quotient formation into a distributed setting. Local state spaces, restricted equivalence relations, overlap compatibility, generated global relations, descent defects, and transformation-stable global quotients are treated as parts of a single algebraic and categorical structure. The resulting framework provides a first rigorous account of how local observations can form global objects and why apparently compatible local identities can nevertheless fail to descend.

1.1 The Stable-Quotient Foundation

The starting point is the algebra of object formation developed for a single state space.

Let X be a state space containing the variations available to a theory, observer, instrument, or computational system. An observation policy determines which distinctions are retained. It can be represented canonically by an equivalence relation E on X.

If x E y, then x and y count as the same at the selected conceptual resolution.

The quotient

X/E

is the collection of equivalence classes determined by E. Each class is an effective object. The quotient therefore represents the object space made available by the observation policy.

This construction does not imply that the underlying states cease to exist or become literally identical. It states that their differences are irrelevant for the purpose represented by E. Quotient identity is policy-relative identity.

The quotient alone is not sufficient for lawful objecthood. Suppose a transformation g acts on X. For g to induce a well-defined transformation on X/E, it must preserve E-equivalence:

x E y implies g(x) E g(y).

When this condition holds, the transformation descends uniquely to the quotient:

$\tilde{g}([x]E) = [g(x)]E$.

Representative independence follows because equivalent representatives have equivalent images. The quotient object can therefore be transformed without referring to the particular state chosen to represent it.

A quotient satisfying this condition for the relevant family of transformations is a stable quotient. Stable quotienting supplies a formal transition from detailed states to effective objects:

state variation,

observation policy,

equivalence,

quotient formation,

transformation invariance,

stable objecthood.

The family of all equivalence relations on X can be ordered by refinement. A finer policy preserves more distinctions. A coarser policy identifies more states. Strict refinement creates relative object birth by splitting a previously unified class. Coarsening produces object fusion by merging several classes into one.

These policies form a complete lattice. Their intersections produce common refinements, while equivalence closures of unions produce common coarsenings. The refinement order also forms a thin category in which arrows represent the lawful forgetting of distinctions.

Maps between state spaces transport policies backward by pullback. If f maps X into Y and F is an equivalence relation on Y, then the induced relation on X is defined by

x f*F y exactly when f(x) F f(y).

This gives a contravariant organization of observational resolution.

These results establish the single-domain theory. They explain how one policy produces stable objects, how policies compare, and how identity structures move between state spaces.

They do not yet explain how multiple policies defined only on parts of a space can be assembled into one global policy. The present paper addresses that missing construction.

1.2 Why a Single Observation Policy Is Insufficient

The single-policy framework assumes that one equivalence relation is already defined across the relevant state space. This assumption is appropriate when the observer has global access, when the classification system is centrally specified, or when the mathematical model is constructed from the beginning as one unified domain.

Many real systems do not have that form.

Observation may be spatially local. A sensor records only the states within its detection region. A medical imaging device resolves one section of an organ. An astronomical instrument sees one field. A robot perceives only the portion of its environment within range.

Observation may be temporally local. Measurements made at different times may overlap only through persistent markers, trajectories, or inferred continuities.

Observation may be institutionally local. Different laboratories, agencies, databases, or disciplines may maintain independent identity systems. Their categories may agree in some cases and diverge in others.

Observation may be representationally local. One model describes geometric features, another causal behavior, another statistical similarity, and another functional role. Each recognizes objects through different coordinates or attributes.

Observation may be computationally local. Specialized models classify different subsets of a large dataset. Distributed agents possess only partial state information. Privacy or communication constraints prevent the construction of one complete centralized representation.

Observation may also be intrinsically contextual. The distinction structure available in one experimental or theoretical context may not be simultaneously available in another. There may be no single observational standpoint from which all relevant identities can be directly compared.

In such settings, a global policy is not given. It must be constructed.

Let X be a state space covered by subsets Ui. Each Ui represents the domain accessible to observer i. On each Ui, an equivalence relation Ei determines local identity. The family

{Ei on Ui}

is the distributed observation system.

The central question is whether there exists an equivalence relation E on X such that

E restricted to Ui equals Ei

for every i.

If such an E exists, the local policies are the exact restrictions of one global identity policy. The local objects can then be interpreted as partial views of a coherent global object system.

If no such E exists, the distributed observations cannot be combined without changing at least one local identity judgment.

This issue cannot be reduced to ordinary stable quotienting. Stability asks whether transformations preserve an existing equivalence relation. Distributed object formation asks whether the equivalence relation itself can be globally formed from local data.

The problem of object stability must therefore be preceded by the problem of object descent.

1.3 Local Knowledge and Partial Object Formation

A local observer does not see a fragment of a fully specified global quotient in any automatic sense. The observer sees a domain Ui equipped with a local identity policy Ei. Its object space is the quotient

Ui/Ei.

These quotient classes are locally meaningful objects. They may support local prediction, classification, and transformation even when no global object system has been established.

A local object can be complete relative to the observer's domain. A physician may identify a pathological structure in one image slice. A field biologist may classify organisms within one ecosystem. A camera may track an object through one sector. A mathematical chart may represent a manifold correctly within its coordinate neighborhood.

The word local therefore does not mean inaccurate or incomplete in every respect. It means that the scope over which identity has been certified is restricted.

Difficulties arise when local object classes cross domain boundaries.

Suppose a class in Ui intersects Ui ∩ Uj. The overlap permits comparison with the classes of Ej. If both policies agree on the overlap, the shared states receive consistent local identities. The overlap then acts as an interface through which local objects may be related.

A network of overlaps can propagate identity across the cover. A class in U1 may connect to a class in U2, which connects to a class in U3. Through such chains, a candidate global object may be assembled.

The resulting global object is not necessarily present as one class in any local quotient. It may exist only as a connected family of compatible local classes. Distributed objecthood can therefore exceed the representational capacity of every individual observer.

This is one reason local-to-global construction is mathematically important. The whole is not obtained merely by collecting the local objects as independent pieces. Identity can travel through overlaps and create a larger unit.

At the same time, that propagation can become excessive. A chain may return to a previous domain and identify two classes that the local observer had intentionally separated. The network has then generated an identity that conflicts with the original local resolution.

Local knowledge therefore has two roles.

It supplies the evidence from which global objects may be assembled.

It also supplies the constraints that any faithful globalization must preserve.

A valid global theory must respect both roles.

1.4 The Local-to-Global Problem

The local-to-global problem can be stated precisely.

Let X be a state space and let {Ui} be a cover of X:

$X = \bigcup_i U_i$.

For each i, let Ei be an equivalence relation on Ui.

We seek a global equivalence relation E on X such that

$E|U_i = E_i$

for every i.

The first necessary condition is agreement on overlaps. For every pair i and j,

$E_i|U_i \cap U_j = E_j|U_i \cap U_j$.

If two local observers disagree about the identity of states they both observe, no global policy can restrict exactly to both of their relations.

Overlap agreement is therefore necessary.

It is not sufficient.

The problem is that a global equivalence relation must account for chains crossing several domains. Suppose

$x \; E_1 \; y$,

$y \; E_2 \; z$,

$z \; E_3 \; w$.

Any global equivalence containing the local relations must identify x with w by transitivity. If x and w also lie together in one local domain Uk but are distinguished by Ek, then faithful globalization is impossible.

No pairwise overlap need contain the entire chain. Each local agreement may be perfectly consistent. The contradiction appears only after the identifications are composed globally.

This reveals the true structure of descent. Local data must satisfy not only direct compatibility but closure compatibility. They must remain unchanged under all global equivalence consequences generated by the network.

The local-to-global problem therefore has three layers:

local equivalence within each domain,

compatibility on overlaps,

global closure without local distortion.

Only when all three are satisfied do the local objects descend faithfully.

1.5 Agreement on Overlaps Versus Global Coherence

Agreement on overlaps is intuitive because overlaps are the regions in which local observers can directly compare their classifications. If two policies assign the same equivalence relation to every shared region, they appear to be mutually consistent.

For many mathematical structures, compatible local data can be glued uniquely into a global structure. Functions, sections of suitable bundles, and other sheaf-like objects are governed by familiar local agreement conditions. This success can encourage the assumption that equivalence relations behave similarly.

They do not behave like ordinary pointwise data.

An equivalence relation contains a global closure law. Transitivity can connect endpoints through intermediate states lying outside their shared domain. Consequently, the value of a relation on one subset can be altered by identity paths that leave the subset and later return.

Suppose U and V overlap. A relation on U identifies a with b. A relation on V identifies b with c. If a and c both lie in U but the relation on U distinguishes them, any global equivalence containing both local relations contradicts the original policy on U.

The local relations may agree perfectly on $U \cap V$. The failure is caused not by disagreement on shared pairs, but by a global path that exits U through b and returns through c.

This phenomenon may be called identity leakage. An identification originating outside a domain reenters it and collapses distinctions that were locally preserved.

Global coherence therefore requires more than agreement on visible interfaces. It requires immunity to all identification paths generated by the entire cover.

A distributed observation network is globally coherent when every global zig-zag between two states already respects the local policy of any domain containing both endpoints.

This condition is stronger than pairwise compatibility and captures the essential obstruction to faithful descent.

1.6 Transitivity as the Hidden Gluing Problem

Reflexivity and symmetry usually cause little difficulty in globalization. Every state remains identical to itself, and local symmetric relations remain symmetric when combined.

Transitivity is the decisive operation.

Let R be the union of the local equivalence relations:

$R = \bigcup_i E_i$.

The relation R contains every identity explicitly licensed by at least one local observer. It is reflexive if the cover includes every state and each Ei is reflexive on its domain. It is symmetric because every Ei is symmetric.

But R need not be transitive.

There may exist states x, y, and z such that

x R y,

y R z,

but not x R z.

To obtain a global equivalence relation, one must add all identities required by transitivity. Repeating this process produces the equivalence closure

$E^* = \mathrm{EqCl}(R)$.

The relation E* is the smallest equivalence relation on X containing all local policies. It is therefore the canonical minimal globalization.

Every global equivalence relation that accepts all local identifications must contain E*.

This universal property makes E* unavoidable. If E* introduces an unwanted local identification, no alternative global equivalence can avoid that identification while still containing all the local relations.

The obstruction is therefore structural rather than an artifact of a poor gluing algorithm.

Transitivity generates paths of identity. A finite zig-zag

$x = x_0, x_1, \ldots, x_n = y$

with each consecutive pair identified in some local domain forces x E* y.

The global equivalence classes are the connected components of this identity network.

This graph-like interpretation is useful, but it must be handled carefully. Connectivity creates the global classes, while faithful descent requires that connectivity not produce extra connections within any one local domain.

The hidden gluing problem can therefore be expressed as follows:

Does transitive closure preserve every local equivalence relation exactly?

When the answer is yes, the local policies descend.

When the answer is no, a gluing obstruction exists.

1.7 Principal Contributions and Scope

This paper develops a distributed extension of the stable-quotient framework through eight principal contributions.

First, it defines local observation systems as covers equipped with equivalence relations on their domains. Each local relation produces a quotient object space representing the objects visible to that observer.

Second, it formalizes overlap compatibility by requiring local policies to have identical restrictions on shared domains.

Third, it constructs the minimal globalization of a compatible or incompatible family as the equivalence closure of the union of all local relations.

Fourth, it proves a faithful descent criterion. The local family descends exactly when restricting the generated global equivalence to each local domain recovers the original local policy.

Fifth, it defines gluing defects as the additional local identifications introduced by global closure. These defects can be represented by explicit zig-zag witnesses and quantified in finite systems.

Sixth, it distinguishes faithful gluing, failed gluing, false gluing, and emergent global identity. This separates direct inconsistency from excessive closure and distinguishes destructive fusion from legitimate network-generated objecthood.

Seventh, it extends transformation stability to the distributed setting. Compatible local transformations can induce a lawful action on the global quotient when they preserve the generated equivalence and assemble into a coherent global transformation.

Eighth, it studies the assignment of resolution lattices to observational domains as a presheaf-like structure. Restriction is contravariant, but faithful gluing can fail. The family of all equivalence relations therefore exposes a precise limit of naive sheaf behavior.

The scope remains deliberately foundational.

The paper treats exact equivalence relations on sets. It does not yet develop probabilistic, fuzzy, graded, metric, or homotopical identity.

It assumes that observational domains are represented by subsets or by similarly well-behaved local state spaces. More general sites, topologies, and indexing categories are deferred.

It does not claim that every distributed scientific object is fully represented by an equivalence class. The theory isolates one exact and widely applicable mechanism of object assembly.

It does not resolve automatically whether a newly generated global identity is valid emergence or destructive fusion. That judgment requires a specified transformation regime, task, explanatory objective, or admissibility policy.

It also does not yet construct a full stack or higher category of observation procedures. Such structures become relevant when local policies are equivalent rather than literally equal, when multiple gluing witnesses exist, or when the compatibility data themselves possess nontrivial transformations.

The present goal is to establish the first-order descent theory clearly enough that these higher extensions rest on a rigorous base.

1.8 Main Results

The main results can be summarized through one local-to-global construction.

Let X be a state space covered by a family {Ui}. For each i, let Ei be an equivalence relation on Ui.

Define the local identity relation R on X by

R = $\bigcup_i$ Ei.

Define the generated global relation by

E* = EqCl(R).

Then E* is the smallest equivalence relation on X containing every local policy Ei.

This is the minimal globalization theorem.

A faithful global equivalence relation exists precisely when

$E^*|U_i = E_i$

for every i.

When this condition holds, E^* is the unique smallest faithful globalization. Every other global equivalence relation containing all local policies must contain E^*. If another global relation restricts exactly to every E_i, it may make additional identifications only between states that never occur together in a domain whose local policy would distinguish them.

The failure of faithful descent on U_i is measured by the defect relation

$D_i = E^*|U_i \setminus E_i$.

A pair belongs to D_i when it is locally distinguished but globally identified through a zig-zag of local relations.

The family descends faithfully exactly when

$D_i = \emptyset$

for every i.

Each nonempty defect admits a finite witness under the ordinary finitary definition of equivalence closure: a chain of locally licensed identifications joining two locally distinguished endpoints.

If transformations act locally, additional compatibility conditions are required. Local transformations must agree on overlaps so that they assemble into a global map. They must also preserve the local relations. Under suitable assembly and invariance conditions, the resulting global transformation preserves E^* and descends to the global quotient X/E^*.

Thus distributed stable objecthood requires two forms of coherence:

descent coherence of identity policies,

dynamical coherence of admissible transformations.

The categorical structure follows from restriction. Each domain U is assigned its lattice Eq(U) of observation policies. An inclusion $V \subseteq U$ induces a restriction map

$Eq(U) \to Eq(V)$.

These assignments are contravariant and compositional. They form a presheaf of conceptual resolutions. However, compatible local relations need not possess a faithful global gluing. The presheaf therefore fails the unrestricted sheaf condition in a mathematically informative way.

The failure is not merely negative. It identifies the exact place where global object formation adds structure beyond local observation. Equivalence closure can create identities that are invisible locally. The resulting gluing obstruction is therefore also a measure of potential emergence.

The combined result establishes a local-to-global algebra of object formation:

local observation produces local identity,

overlap comparison produces compatibility data,

equivalence closure produces minimal globalization,

restriction tests faithful descent,

defect relations expose false gluing,

transformation invariance certifies global object stability.

The theory therefore extends mathematical object formation from the question

How does one observation policy produce stable objects?

to the more general question

How can many partial policies jointly produce one coherent world of objects?

2. Local State Spaces and Restricted Observation Policies

The local-to-global theory begins with a distinction between the total state space and the domains accessible to particular observers. A global state space contains all states relevant to the system under study, while a local observational domain contains only those states that can be compared, measured, classified, or transformed within one restricted context.

This distinction is structural rather than merely practical. A local domain may be restricted because of spatial range, temporal access, experimental context, computational partitioning, institutional separation, representational choice, or theoretical scale. The observer does not initially possess a global equivalence relation. It possesses an equivalence relation only on the states jointly available within its domain.

Let X be a state space.

Let U be a subset of X.

An observation policy on U is an equivalence relation E on U.

The pair

(U, E)

is a local policy-indexed state space.

Its quotient

U/E

is the corresponding local object space.

Each element of U/E is an equivalence class of states that count as one object under the local policy. The scope of that identity claim is U. Nothing in the local relation alone determines whether a class extends beyond U, corresponds to part of a larger object, or remains isolated when compared with other observational domains.

A distributed observation system consists of a family

$\{(U_i, E_i)\}_{i \in I}$

such that the domains U_i cover X or cover the portion of X relevant to the problem.

The central local data are therefore not merely subsets and not merely equivalence relations. They are equivalence relations indexed by domains. Identity is always stated together with the region over which it has been certified.

This section develops the elementary structure of those local systems. It defines observational domains, covers, restricted policies, local quotients, refinement, and restriction maps. These constructions provide the presheaf-like base on which the later descent theory will operate.

2.1 Global State Spaces and Observational Domains

A state space is a collection of possible states, configurations, descriptions, measurements, or presentations. Its interpretation depends on the application.

In a physical system, X may be a phase space, configuration space, or collection of experimentally accessible states.

In image recognition, X may be a dataset of images or a space of possible visual inputs.

In biology, X may be a collection of organisms, genomes, phenotypes, tissue states, or ecological configurations.

In computation, X may be a set of program states, memory configurations, execution traces, or input-output behaviors.

In scientific knowledge integration, X may be a collection of records, observations, entities, or propositions.

In mathematics, X may be a set, space, category of presentations, collection of models, or domain of candidate structures.

A local observational domain U is a subset of X whose elements can be jointly compared under one observation policy.

The word jointly is important. A policy of identity requires more than isolated access to individual states. It requires the ability to determine whether pairs of states count as the same. The domain U therefore represents the range within which an observer can make coherent identity judgments.

Different domains may arise for different reasons.

A spatial domain may consist of states located within one sensor's range.

A temporal domain may consist of states observed during one time interval.

A laboratory domain may consist of samples processed under one protocol.

A coordinate domain may consist of points represented within one chart.

A computational domain may consist of data stored on one machine.

A disciplinary domain may consist of entities recognized within one scientific vocabulary.

A contextual domain may consist of states jointly measurable under one experimental arrangement.

A scale domain may consist of states visible at one level of resolution.

The same subset of X may support several different observation policies. A collection of biological specimens may be classified by morphology, genotype, ecological role, or developmental lineage. A set of images may be grouped by depicted object, style, source, or semantic content. A set of physical states may be grouped by energy, symmetry orbit, thermodynamic variables, or asymptotic behavior.

Thus the domain determines where identity is evaluated, while the policy determines what identity means there.

The pair (U, E) should therefore be treated as the primitive local observational object.

The underlying subset U alone does not specify which states form objects.

The relation E alone does not specify the scope within which its identity judgments are meaningful.

Together they define a local object-forming context.

A local domain need not be maximal. It may be enlarged if additional states become observable, or restricted if the observer loses access or if a smaller context is required. Changing the domain can alter which equivalence relations are available, which comparisons can be made, and which objects can be formed.

Domain variation is therefore one component of conceptual change.

Policy variation is another.

The local-to-global theory must track both.

2.2 Covers by Local Regions

A family of subsets $\{U_i\}_{i \in I}$ covers X when every state belongs to at least one observational domain:

$X = \bigcup_{i \in I} U_i$.

The index set I labels observers, instruments, contexts, coordinate patches, computational agents, laboratories, scales, or other local viewpoints.

The cover condition ensures that no state in X is completely absent from the distributed observation system. It does not require every pair of states to occur together in one domain. That stronger property would make many local-to-global difficulties disappear, but it is rarely realistic.

A cover may be disjoint.

In that case, no two distinct domains share states, and there is no direct observational interface between them. Local policies can be combined without overlap conflict, but global identity between components remains underdetermined unless additional bridging data are introduced.

A cover may overlap sparsely.

Each domain may share states with only a few neighbors. This structure is common in sensor networks, spatial atlases, distributed databases, and coordinate systems.

A cover may overlap densely.

Many observers may see the same states, producing redundancy, cross-validation, or conflicting identity judgments.

The pattern of overlaps can be represented by a nerve. The vertices correspond to domains. A collection of vertices spans a simplex when the corresponding domains have nonempty common intersection. This combinatorial object records which observers can compare their policies directly.

Although the full theory of nerves and Čech constructions is deferred, the overlap pattern already matters at the elementary level. It determines which local relations can be compared and which chains of identity may propagate across the network.

A cover may also be refined.

A cover $\{V_a\}_{a\in A}$ refines $\{U_i\}_{i\in I}$ when every V_a is contained in at least one U_i.

Cover refinement increases localization. It replaces larger observational domains with smaller ones. This can reveal incompatibilities that were hidden at coarse scale, but it can also remove some comparisons and weaken global constraints.

Conversely, enlarging domains increases the number of pairs that one observer can compare directly. This can prevent false gluing by making local distinctions explicit, but it may also introduce direct conflicts among formerly separated policies.

The geometry of the cover is therefore not neutral. It influences which global identities can be certified and which gluing defects can be detected.

For the basic theory, no topology on X is required. A cover may be any family of subsets whose union is X. Later applications may impose additional structure, such as open covers, measurable domains, subcomplexes, graph neighborhoods, or objects of a Grothendieck topology.

The set-theoretic theory is deliberately general. It isolates the algebra of local identity before adding geometric or topological assumptions.

2.3 Observation Policies on Subsets

Let U be a subset of X.

An observation policy on U is an equivalence relation $E \subseteq U \times U$.

For $x, y \in U$, the statement

$x \; E \; y$

means that x and y count as the same object within the context represented by U and E.

The relation E must satisfy three conditions.

For every $x \in U$,

x E x.

For every x, y ∈ U,

x E y implies y E x.

For every x, y, z ∈ U,

x E y and y E z imply x E z.

These are reflexivity, symmetry, and transitivity.

Each condition has a local interpretation.

Reflexivity states that every locally available state belongs to some local object.

Symmetry states that local indistinguishability does not depend on the order of comparison.

Transitivity states that local identity is coherent across chains of comparison within the domain.

The equivalence classes of E partition U. Every state belongs to exactly one class, and distinct classes do not overlap.

An observation policy may arise directly from an observation map.

Let

o: U → O

assign an observational output to each state.

The kernel relation of o is defined by

x ker(o) y exactly when o(x) = o(y).

This kernel is an equivalence relation. States belong to the same local object when they produce the same local observation.

Different observation maps may define the same policy. A relabeling of outputs changes the presentation but not the kernel. For object formation, the equivalence relation is therefore more fundamental than the particular labels assigned to its classes.

A local policy may also be defined through invariants, orbits, behavior, or admissible substitution.

Two states may be equivalent because they share the same measured features.

They may be equivalent because a symmetry transformation connects them.

They may be equivalent because they produce the same future behavior.

They may be equivalent because every local test gives the same result.

They may be equivalent because one can replace the other in every locally admissible operation.

These interpretations differ in origin but produce the same formal structure once represented by an equivalence relation.

A domain U may carry the equality relation ΔU:

$x \Delta U y$ exactly when $x = y$.

This is the finest local observation policy. It preserves every available distinction.

The domain may also carry the universal relation ∇U:

$x \nabla U y$ for all $x, y \in U$.

This is the coarsest local observation policy. It identifies every state as one object.

Most useful policies lie between these extremes.

A local observer preserves enough distinctions to support the task while erasing enough variation to form manageable objects.

2.4 Restriction of Equivalence Relations

Suppose $V \subseteq U$ and E is an equivalence relation on U.

The restriction of E to V is

$E|V = E \cap (V \times V)$.

Equivalently, for $x, y \in V$,

$x \, E|V \, y$ exactly when $x \, E \, y$.

The restricted relation is an equivalence relation on V.

Reflexivity is preserved because every $x \in V$ satisfies $x \, E \, x$.

Symmetry is preserved because $x \, E \, y$ implies $y \, E \, x$.

Transitivity is preserved because $x \, E \, y$ and $y \, E \, z$ imply $x \, E \, z$ for all $x, y, z \in V$.

Restriction therefore transports an observation policy from a larger domain to a smaller one without changing any identity judgment among states that remain available.

The local quotient $V/(E|V)$ need not be a subset of U/E in a literal set-theoretic sense, but it maps naturally into the collection of global E-classes that intersect V. Each restricted class has the form

$[x]E \cap V$.

A single E-class may intersect V partially. Restriction does not split that intersection further, but it may remove all representatives connecting the visible portion to states outside V.

This distinction matters epistemically. An observer on V may see a complete local class without knowing that it is part of a larger E-class extending beyond V.

Restriction forgets states, not identity judgments among retained states.

The assignment

$E \mapsto E|V$

defines a map

$res_{U,V}: Eq(U) \to Eq(V)$,

where $Eq(U)$ denotes the set of equivalence relations on U.

If $W \subseteq V \subseteq U$, then

$(E|V)|W = E|W$.

Restriction is therefore compositional.

Restricting to U itself changes nothing:

$E|U = E$.

These laws provide the contravariant structure required later for the presheaf of conceptual resolutions.

Restriction also preserves refinement. If E is finer than F on U, meaning

$E \subseteq F$,

then

$E|V \subseteq F|V$.

Thus a distinction preserved globally remains preserved locally. A coarse global policy cannot become finer merely by restriction, although some differences between policies may disappear if the states witnessing them lie outside V.

Restriction need not preserve joins in the lattice of equivalence relations. A global join may contain transitive identifications whose witnessing paths leave V. Restricting after global closure can therefore produce more local identifications than first restricting the policies and then forming their local join.

This failure is not incidental. It is an elementary form of the same transitivity problem that later obstructs descent.

Restriction does preserve arbitrary meets because meets are intersections:

$(\bigcap_\alpha E_\alpha)|V = \bigcap_\alpha (E_\alpha|V)$.

The asymmetry between meets and joins reflects a basic difference.

Combined distinction is local.

Generated identification may be nonlocal.

2.5 Local Quotient Object Spaces

Given a local policy-indexed state space (U, E), the quotient

U/E

is the set of all E-equivalence classes.

The canonical quotient map is

$q_E: U \to U/E$.

It sends each state x to its class $[x]_E$.

The quotient map is surjective. Every local object has at least one representative in U.

The quotient forgets distinctions within each class while preserving distinctions between classes. It therefore represents the object ontology available to the local observer.

A local object should not be confused with a globally existing object fragment in advance of descent. It is an object relative to the domain and policy. Its relation to objects in other domains remains to be established.

Suppose U and V overlap. A local class $A \in U/E$ and a local class $B \in V/F$ may intersect within $U \cap V$. If $A \cap B$ is nonempty, they share at least one state. If the policies agree on the overlap, that shared state may support an identification between the local objects.

However, two local classes may correspond globally even if they do not directly intersect. A chain of overlapping classes across several domains may connect them. Global object formation

therefore occurs not only by matching classes pairwise but by taking connected components of a larger identity network.

The size of U/E measures the number of locally distinguishable objects when U is finite.

The compression factor may be written as

$\kappa(U, E) = |U| / |U/E|$.

A large value indicates that many local states have been compressed into comparatively few objects. Cardinal compression alone does not establish the quality of the policy. Stability and descent remain necessary.

Two local policies may have the same number of quotient classes while preserving entirely different distinctions. Object counts do not determine object identity.

A quotient also inherits transformations when they preserve E. If $g: U \to U$ satisfies

x E y implies g(x) E g(y),

then g induces

$\bar{g}: U/E \to U/E$.

Thus local object spaces can support lawful local dynamics before any global system exists.

This is important in distributed applications. Each observer may reason correctly within its own quotient while the collection of local quotients fails to combine coherently.

Local correctness does not imply global coherence.

2.6 Refinement Within Observational Domains

Let E and F be observation policies on the same domain U.

E is finer than F when

$E \subseteq F$.

This means that every pair identified by E is also identified by F, but F may identify additional pairs.

The finer policy preserves at least as many distinctions as the coarser policy.

Write

$E \preceq F$

for the refinement order.

Under this convention, equality is the bottom element and the universal relation is the top element.

The quotient U/E contains at least as many classes as U/F in the finite case. Each F-class is a union of E-classes.

There is a canonical surjection

$\pi_{E,F}: U/E \to U/F$

defined by

$\pi_{E,F}([x]_E) = [x]_F$.

This map is well defined because $x \, E \, y$ implies $x \, F \, y$.

The canonical surjection represents lawful forgetting. A fine object is sent to the coarser object containing it.

Strict refinement occurs when

$E \subset F$.

At least one F-class then splits into multiple E-classes. Relative to F, the finer policy produces object birth.

Strict coarsening merges previously distinct objects. Relative to E, the coarser policy produces object fusion.

Refinement is domain-relative. Two policies can be compared directly only when they are defined on the same domain, or after both have been restricted to a common subdomain.

This observation becomes important on overlaps. Policies E_i on U_i and E_j on U_j are compared by restricting both to $U_i \cap U_j$.

Exact agreement requires equality of restrictions.

One may also study asymmetric agreement. The restriction of one policy may refine the other, indicating that one observer preserves distinctions the other erases. Such relations will be relevant to approximate integration and hierarchical observers, although the present paper focuses first on exact compatibility.

The collection Eq(U), ordered by refinement, forms a complete lattice.

For a family $\{E_\alpha\}$, the meet is the intersection

$\bigwedge_\alpha E_\alpha = \bigcap_\alpha E_\alpha$.

The meet preserves every distinction preserved by any policy in the family.

The join is the equivalence closure of the union:

$\bigvee_\alpha E_\alpha = \mathrm{EqCl}(\bigcup_\alpha E_\alpha)$.

The join is the smallest policy coarser than every E_α.

The same transitive closure that creates joins inside Eq(U) later generates global relations across domains. The local lattice theory and the global descent theory are therefore two instances of one operation: the completion of declared identities under equivalence.

2.7 Restriction Maps Between Resolution Lattices

For every inclusion $V \subseteq U$, restriction defines a monotone map

$\mathrm{res}_{U,V}: \mathrm{Eq}(U) \to \mathrm{Eq}(V)$.

These maps satisfy

$\mathrm{res}_{U,U} = \mathrm{id}_{\mathrm{Eq}(U)}$

and

$\mathrm{res}_{V,W} \circ \mathrm{res}_{U,V} = \mathrm{res}_{U,W}$

whenever $W \subseteq V \subseteq U$.

Thus the assignment

$U \mapsto \mathrm{Eq}(U)$

reverses inclusions.

A larger domain maps to a smaller domain by restriction of policy.

This is contravariance.

At the order-theoretic level, each Eq(U) is a complete lattice, and each restriction map preserves arbitrary meets:

$\mathrm{res}_{U,V}(\bigwedge_\alpha E_\alpha) = \bigwedge_\alpha \mathrm{res}_{U,V}(E_\alpha)$.

Restriction also preserves the bottom and top policies:

$\Delta_U|_V = \Delta_V$,

$\nabla_U|_V = \nabla_V$.

It does not generally preserve arbitrary joins.

To see why, suppose E and F are relations on U. Their join may identify x and z through a chain

x E y,

y F z,

where y lies outside V while x and z lie inside V.

Then x and z may be related in

$(E \vee F)|V$

even though they are not related in

$E|V \vee F|V$.

Therefore,

$(E \vee F)|V$

may be strictly coarser than

$E|V \vee F|V$.

This inequality expresses identity leakage through restriction. A global path leaves V, creates an identification, and returns.

Restriction forgets the witness but retains the consequence.

This phenomenon anticipates the main gluing obstruction. A relation generated globally can impose identities inside a local domain that cannot be generated from the local data alone.

The family Eq(U) should therefore not be regarded as a sheaf of lattices with unrestricted gluing properties. It is more naturally a contravariant system of complete meet-semilattices with additional nonlocal behavior in its joins.

The exact categorical formulation will be developed later. At this stage, the important fact is that local observation policies restrict coherently, while coarsening by generated equivalence may depend on states outside the restricted domain.

This local-global asymmetry is the algebraic source of descent failure.

2.8 Examples of Locally Formed Objects

Several examples illustrate why local policies must be indexed by domains and why local object formation does not automatically determine global identity.

Consider a network of cameras observing adjacent sectors of a building. Each camera groups image detections into local tracks. Within one sector, appearances at successive frames are identified as one moving object. The quotient consists of local tracks rather than individual images.

On overlapping sectors, two cameras may observe the same person. Agreement permits their local tracks to be connected. A chain of track connections can produce one global trajectory.

False gluing occurs if two visually similar people cross different sectors and the network joins their tracks through ambiguous intermediate detections. Each camera may have classified its own observations consistently, yet the global closure may fuse distinct people.

Consider an anatomical atlas constructed from image slices. Each slice is a domain. A local policy groups pixels or voxels into tissue structures. Adjacent slices overlap through registration or shared landmarks.

A three-dimensional organ is not present as a complete object in any one slice. It emerges from the gluing of local classes. Incorrect correspondence between slices can produce a false anatomical fusion, while failure to match corresponding regions can fragment one organ into multiple global objects.

Consider biological classification across geographic regions. Each research group studies organisms within one region and defines local species or population classes. Overlap zones permit comparison.

A chain of interbreeding populations may connect distant groups that do not directly interbreed. Depending on the chosen identity policy, transitive closure may produce a ring-species-like global class or may erase biologically significant distinctions. The mathematics identifies the globalization; biology must determine whether the resulting identity is appropriate.

Consider a distributed database. Each institution maintains records and defines when two records refer to the same entity. Shared identifiers create overlap relations.

Entity resolution across databases requires transitive closure. If record A matches record B and record B matches record C, the system may merge all three. An erroneous bridge record can cause a large cluster of false identities. The local records may contain no direct contradiction, yet the global entity class becomes destructive.

Consider coordinate charts on a geometric space. Each chart provides a local representation of points. Transition functions identify points on overlaps.

When the transition data satisfy the required compatibility conditions, the local descriptions form a coherent global space. This familiar geometric case motivates descent theory, although the present framework differs because local equivalence relations can generate additional identities through closure.

Consider specialized artificial intelligence models. One model classifies medical images, another analyzes clinical text, and another processes genomic data. Each operates on a different feature domain and forms local patient or disease objects.

Shared cases provide overlaps. Integration may reveal a global disease subtype not represented within any single model. It may also falsely fuse distinct conditions if transitive associations propagate through weak or ambiguous links.

Consider multiscale physics. A microscopic domain distinguishes detailed particle configurations. A mesoscopic domain groups them into collective modes. A macroscopic domain identifies thermodynamic states.

The domains may overlap through coarse-graining maps rather than literal subset inclusion. Even in this generalized setting, the same questions arise: which identities are preserved, which are induced, and whether local quotient objects form one coherent hierarchy.

These examples share a common structure.

There are restricted domains of comparison.

Each domain supports a local policy of sameness.

Local classes are meaningful objects within their contexts.

Overlaps provide channels for identity propagation.

Global closure may create coherent large-scale objects, fragmentation, contradiction, or false fusion.

The purpose of the theory is not to decide every application-specific identity judgment in advance. It is to make the structure of those judgments explicit and to determine when local object systems can be assembled without changing what the local observers claimed.

The elementary data are now in place:

a state space X,

a cover {Ui},

local equivalence relations Ei,

local quotients Ui/Ei,

restriction maps between domains,

and refinement lattices Eq(Ui).

The next task is to formalize what it means for these local policies to agree on overlaps and to determine why such agreement does not by itself guarantee faithful global objecthood.

3. Overlap Compatibility and Descent Data

Local observation policies become candidates for global object formation only when they can be compared on shared domains. If two observers have no states in common, their identity judgments cannot be tested directly against one another. Where their domains overlap, however, both policies apply to the same states. The overlap becomes the region in which consistency can be evaluated.

Let X be a state space covered by subsets $\{U_i\}_{i \in I}$.

For each i, let Ei be an equivalence relation on Ui.

For every pair of indices i and j, define the overlap

$U_{ij} = U_i \cap U_j$.

The restrictions

$E_i|U_{ij}$

and

$E_j|U_{ij}$

are equivalence relations on the same set Uij. They can therefore be compared directly.

The most immediate compatibility condition is equality:

$E_i|U_{ij} = E_j|U_{ij}$.

This condition states that whenever two states are jointly visible to observers i and j, both observers make the same identity judgment about them.

If x and y lie in Uij, then

x Ei y exactly when x Ej y.

Pairwise overlap agreement is the natural first requirement for distributed observation. A global equivalence relation whose restrictions are the Ei must necessarily satisfy it. If two local policies disagree on an overlap, no global relation can reproduce both exactly.

Yet overlap agreement is only the first layer of descent.

A family of equivalence relations can agree on every pairwise overlap while still failing to arise as the exact restrictions of any global equivalence relation. The reason is that equivalence is not determined point by point or pair by pair in isolation. Transitivity composes local identifications across chains of domains. Those chains may generate new relations between states that lie together in a local domain but were locally distinguished there.

The descent problem must therefore distinguish three ideas:

local validity,

overlap compatibility,

global realizability.

Local validity means that each Ei is an equivalence relation on Ui.

Overlap compatibility means that the restrictions agree wherever domains intersect.

Global realizability means that there exists an equivalence relation E on X whose restriction to every Ui is exactly Ei.

The first two conditions do not automatically imply the third.

This section formalizes the data available on overlaps, explains why pairwise agreement is necessary but insufficient, and identifies the additional closure condition required for faithful descent. The term descent data will refer to a compatible family of local identity policies together with the information required to determine whether those policies assemble into a coherent global object system.

3.1 Pairwise Intersections of Observational Domains

For two observational domains Ui and Uj, the intersection

$U_{ij} = U_i \cap U_j$

contains the states accessible to both local observers.

The overlap may be empty. In that case, the two policies cannot be compared directly. Their local objects remain observationally disconnected unless other domains form a bridge between them.

If Uij is nonempty, both Ei and Ej restrict to equivalence relations on Uij.

The overlap has two related roles.

First, it is a comparison domain. The observers can test whether they agree about the identity of shared states.

Second, it is a transmission domain. Shared states can connect equivalence classes from Ui to equivalence classes from Uj, allowing object identity to propagate across the cover.

Suppose x belongs to Uij. The local class of x under Ei is

$[x]E_i \subseteq U_i$.

The local class of x under Ej is

$[x]E_j \subseteq U_j$.

These classes need not be equal because they live in different domains. They share at least the state x, and they may contain different additional states visible only to one observer.

If the policies agree on Uij, then

$[x]E_i \cap U_{ij} = [x]E_j \cap U_{ij}$.

The two local classes have the same visible trace on the overlap.

This does not mean that the classes are already one global object. It means only that their shared portions are identified consistently.

An overlap may contain several local classes. Agreement requires that both policies induce exactly the same partition of the overlap. Each shared state belongs to the same overlap class under both observers.

Thus compatibility can be expressed either relationally,

$E_i|U_{ij} = E_j|U_{ij}$,

or partition-theoretically:

the partitions induced by Ei and Ej on Uij are identical.

The relational formulation is generally preferable because it interacts directly with restriction, refinement, closure, and categorical structure.

The size and shape of overlaps influence the strength of compatibility testing.

A one-point overlap can connect local objects but cannot test any nontrivial pairwise identity judgment. Every equivalence relation restricts to the same relation on a singleton.

A larger overlap can reveal whether local classes align.

A richly connected system of overlaps may strongly constrain global identity.

A sparse system may leave many global identifications underdetermined.

Overlap geometry therefore affects not only data sharing but also the amount of object structure that can be certified.

3.2 Agreement of Identity Policies on Overlaps

The family {Ei} is pairwise compatible when, for all i and j,

Ei|Uij = Ej|Uij.

This equality can be unpacked as follows.

For every x, y ∈ Ui ∩ Uj,

x Ei y if and only if x Ej y.

Pairwise compatibility means that no shared pair is declared identical by one observer and distinct by another.

This condition is necessary for exact global descent.

Suppose E is a global equivalence relation on X satisfying

E|Ui = Ei

for every i.

Then on Uij,

Ei|Uij = E|Uij = Ej|Uij.

Therefore every globally realizable family is pairwise compatible.

The contrapositive is useful.

If there exist i, j and x, y ∈ Uij such that

x Ei y

but

x is not Ej-equivalent to y,

then no global equivalence relation can restrict exactly to both Ei and Ej.

This is a direct overlap conflict.

Such a conflict is locally visible. It does not require transitive closure or examination of the entire network. The two observers disagree about a pair they both observe.

Direct overlap conflict should be distinguished from descent obstruction.

A direct conflict means the family is not even compatible.

A descent obstruction means the family is pairwise compatible but still cannot be globalized faithfully.

The distinction is important in applications.

In a database, direct conflict occurs when two institutions share the same pair of records but one declares them to represent one entity while the other declares them distinct.

A descent obstruction occurs when a chain of matches through other databases forces those two records to merge globally even though the local database distinguishes them.

In sensor fusion, direct conflict occurs when two sensors jointly observe two detections but disagree about whether they belong to one track.

A descent obstruction occurs when track associations across multiple sensors create a global loop that merges locally separate tracks.

The first problem is disagreement.

The second is closure.

Pairwise compatibility removes disagreement but does not control all consequences of closure.

3.3 Compatible Families of Local Equivalences

A compatible local equivalence system on a cover $\{U_i\}$ consists of a family

$E = \{E_i \in Eq(U_i)\}_{i \in I}$

such that

$E_i|_{U_i \cap U_j} = E_j|_{U_i \cap U_j}$

for every i and j.

This family can be viewed as a section-like assignment of identity policies to the cover.

Each Ei is locally complete as an equivalence relation.

The compatibility equations ensure that the family behaves consistently under restriction to pairwise intersections.

Several elementary consequences follow.

First, compatibility is symmetric in i and j.

Second, compatibility is automatically satisfied when $U_i \cap U_j$ is empty.

Third, compatibility is preserved under common refinement. If each E_i is replaced by a finer policy F_i in a way that preserves equality on overlaps, then the new family is also compatible.

Fourth, compatibility is preserved under common coarsening only when the coarsening procedure itself respects overlaps. Independent local joins can create different induced relations on shared regions.

Fifth, the restriction of a compatible family to a smaller subcover remains compatible.

A compatible family determines a well-defined relation on every overlap independent of which local policy is used to compute it.

For a finite collection of indices $J = \{i_1, \ldots, i_n\}$, define the multiple overlap

$U_J = U_{i_1} \cap \cdots \cap U_{i_n}$.

Pairwise compatibility implies that all restricted policies agree on U_J:

$E_{i_1}|U_J = \cdots = E_{i_n}|U_J$.

This follows because U_J is contained in every pairwise overlap.

Thus exact relational compatibility on pairwise intersections automatically produces consistency on all higher intersections.

This fact must be interpreted carefully.

Higher-overlap equality is automatic at the level of relations.

Global descent is not.

The obstruction does not arise because three observers assign inconsistent relations to a triple overlap after agreeing pairwise. It arises because identity chains can pass through several domains and generate consequences outside any single overlap.

The difficulty is therefore not a failure of finite-overlap bookkeeping. It is a failure of global closure fidelity.

3.4 Cocycle-Like Consistency on Multiple Overlaps

In ordinary gluing problems, local data are often connected by transition maps satisfying cocycle conditions on triple overlaps. For exact equivalence relations defined on literal subsets of one common state space, the situation is simpler.

There are no separate coordinate copies that must be identified by nontrivial transition functions. Shared states are literally the same elements of X. Compatibility is expressed by equality of restricted relations.

On a triple overlap

$U_{ijk} = U_i \cap U_j \cap U_k$,

pairwise compatibility gives

$E_i|_{U_{ijk}} = E_j|_{U_{ijk}} = E_k|_{U_{ijk}}$.

No additional equation is required to make these three restrictions agree.

Nevertheless, cocycle-like structure reappears at the level of local quotient objects.

Suppose x lies in U_{ijk}. Its local classes

$[x]_{E_i}$,

$[x]_{E_j}$,

and

$[x]_{E_k}$

have equal traces on the triple overlap. The correspondence between the Ui-class and the Uj-class induced through x should agree with the correspondence obtained through k whenever both are defined.

This consistency is automatic when correspondences are induced solely by shared underlying states and exact overlap equality. It becomes nontrivial in generalized settings where:

local state spaces are not literal subsets of one X,

shared states are connected by comparison maps rather than equality,

policies agree only up to isomorphism,

multiple object correspondences exist,

or transformations between policies carry nontrivial higher data.

Those generalized situations motivate stacks and higher descent. The present set-theoretic theory remains at the strict level.

Even here, a cocycle-like phenomenon appears through chains.

Consider local identifications

x0 Ei1 x1,

x1 Ei2 x2,

…,

xn−1 Ein xn.

The sequence composes identity judgments drawn from several domains. A global equivalence relation must treat the endpoints x0 and xn as equivalent.

The chain is therefore an identity cocycle in a broad sense: local equivalences compose around the observational network and impose a global consequence.

If the chain begins and ends within one domain Uk, faithful descent requires

x0 Ek xn.

A closed or returning chain that violates this condition is an obstruction witness.

Thus the important higher consistency law is not merely equality on triple overlaps. It is consistency of every composable zig-zag with the local relations at its endpoints.

This is the transitive descent condition.

3.5 Local Object Correspondences

Local quotient spaces provide an object-level view of compatibility.

For each domain Ui, let

Qi = Ui/Ei.

On an overlap Uij, the common restricted relation determines an overlap quotient

Qij = Uij/(Ei|Uij),

which is equally

Uij/(Ej|Uij).

There are natural maps from Qij into the local quotient systems in the following sense.

Each overlap class $C \subseteq U_{ij}$ lies inside a unique E_i-class in U_i and a unique E_j-class in U_j. Thus C determines a pair

$(A_i, A_j) \in Q_i \times Q_j$.

This relation identifies which local objects have a common observational trace.

It need not define a function from Q_i to Q_j.

A local class in Q_i may fail to intersect U_{ij}.

A local class may intersect the overlap in several disconnected-looking portions that nevertheless belong to one restricted equivalence class because E_i already identifies them.

Several local classes in one quotient cannot share the same overlap state, but distinct local classes may have no overlap evidence at all.

The most natural object-level structure is therefore a correspondence or span:

$Q_i \leftarrow Q_{ij} \rightarrow Q_j$.

The left map sends an overlap class to the U_i-object containing it.

The right map sends it to the U_j-object containing it.

These spans encode how local object systems touch.

A global object candidate can be represented by a connected component in the network formed from all local quotient classes and all overlap correspondences.

This quotient-level construction mirrors equivalence closure on states.

Create one node for every local equivalence class.

Connect two nodes when their classes share a state on an overlap.

The connected components of this class network generate candidate global objects.

However, this object-level picture can conceal a descent defect. A connected component may contain two distinct classes belonging to the same local quotient Q_i. In that case, the globalization would merge two objects that observer i explicitly distinguished.

This yields an object-level faithful descent criterion:

No connected component of the local-class network may contain two distinct classes from the same domain.

More precisely, whenever two Ei-classes become connected through overlap correspondences, they must already be the same Ei-class.

This criterion is equivalent to requiring that the generated global equivalence restrict exactly to Ei.

The state-level and quotient-level views are therefore complementary.

The state-level view emphasizes equivalence closure and zig-zags.

The quotient-level view emphasizes networks of local objects and forbidden mergers.

3.6 Why Pairwise Agreement Appears Sufficient

Pairwise compatibility is psychologically and mathematically tempting because every direct comparison succeeds.

Each observer is internally coherent.

Every shared pair receives the same identity judgment.

Every multiple overlap inherits one common restricted relation.

No explicit contradiction is visible.

Under these conditions, it is natural to imagine constructing the global relation by declaring:

x is globally equivalent to y whenever some local policy identifies them.

This produces the union relation

$R = \bigcup_i E_i$.

Because the local relations agree on overlaps, R appears to be a coherent amalgamation.

The difficulty is that R may not be transitive.

Consider four states a, b, c, and d, with domains

$U1 = \{a, b, d\}$,

$U2 = \{b, c\}$,

$U3 = \{c, d\}$.

Let E1 identify a with b while keeping d separate.

Let E2 identify b with c.

Let E3 identify c with d.

Assume all other relations are those forced by equivalence within each domain.

The overlaps are singletons:

$U1 \cap U2 = \{b\}$,

$U2 \cap U3 = \{c\}$,

$U1 \cap U3 = \{d\}$.

All local policies therefore agree on every overlap automatically.

Yet the local identifications generate the chain

a E1 b,

b E2 c,

c E3 d.

Any global equivalence containing the local relations must identify a with d.

But a and d both lie in U1, where E1 distinguishes them.

Therefore no faithful global equivalence exists.

Nothing failed on an overlap. The failure appeared only after composition across the cover.

This example shows why pairwise agreement is insufficient even when every intersection condition is satisfied exactly.

The cover can hide a long-range contradiction.

The contradiction is not an inconsistency in the declared local relations. It is an inconsistency between those declarations and the global transitivity required of equivalence.

Pairwise compatibility checks every edge of the overlap network.

Faithful descent must check every path.

3.7 The Difference Between Compatibility and Descent

Compatibility and descent answer different questions.

Compatibility asks:

Do local policies agree wherever they are directly compared?

Descent asks:

Are the local policies the exact restrictions of one global policy?

Compatibility is local.

Descent is global.

A family can be classified into three preliminary states.

The family is locally invalid if some Ei is not an equivalence relation.

The family is incompatible if the Ei are equivalence relations but disagree on at least one overlap.

The family is compatible if all overlap restrictions agree.

Among compatible families, a further division is required.

The family is descendable when a global equivalence E exists with

$E|U_i = E_i$

for every i.

The family is obstructed when no such E exists.

Thus compatibility is necessary for descent but not sufficient.

The distinction can be expressed using the generated relation.

Let

$R = \bigcup_i E_i$

and

$E^* = \mathrm{EqCl}(R)$.

Every global equivalence containing all local relations must contain E^*.

If

$E^*|U_i = E_i$

for every i,

then E^* itself is a faithful globalization, and the family descends.

If, for some i,

$E_i \subset E^*|U_i$,

then global closure introduces an additional identity inside Ui. No global equivalence containing all local policies can avoid that identity. Therefore faithful descent is impossible.

This yields a complete criterion:

A family of local equivalence relations descends if and only if its generated global equivalence restricts exactly to every local relation.

Pairwise compatibility is contained within this criterion. If E* restricts to Ei and Ej, then their restrictions to Uij agree. But the criterion also tests every nonlocal identity chain.

Descent therefore includes compatibility plus closure fidelity.

3.8 Descent Data for Distributed Observation

For the present theory, descent data consist of the following components:

a state space X,

a cover $\{U_i\}_{i \in I}$,

an equivalence relation Ei on each Ui,

and exact agreement of the restricted relations on every overlap.

This data set is necessary to pose the faithful globalization problem.

It is not yet a guarantee that descent succeeds.

A compatible local system may therefore be called a descent candidate.

Its minimal globalization is obtained by equivalence closure:

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$.

The candidate is effective when

$E^*|U_i = E_i$

for every i.

It is obstructed when this equality fails.

The difference

$D_i = E^*|U_i \setminus E_i$

records the local pairs identified globally but distinguished locally.

The family {Di} is the descent defect.

At the quotient level, descent data can be represented by:

the local object spaces Qi = Ui/Ei,

the overlap object spaces Qij,

and the spans

Qi ← Qij → Qj.

These spans generate a network of local objects.

A descent defect occurs when one connected component of that network contains multiple distinct objects from the same Qi.

The two formulations are equivalent in the strict set-theoretic setting.

The relation formulation is better for proofs.

The quotient-network formulation is often better for computation and interpretation.

Descent data should not be understood as merely passive records of agreement. They describe a distributed mechanism of object production. Each local observer forms objects independently. Overlaps allow those objects to be related. Closure determines which local objects participate in one global identity. Restriction then tests whether the global result preserves the original local distinctions.

This produces the sequence:

local policy,

local quotient,

overlap comparison,

object correspondence,

identity propagation,

global closure,

descent test.

The distinction between a compatible family and an effective descent family is the central conceptual result of this section.

Compatibility certifies that no direct disagreement is visible.

Descent certifies that no indirect chain creates a forbidden identity.

The next section constructs the generated global equivalence formally, proves its minimality, and establishes its universal role in every possible globalization.

4. Generated Globalization and Equivalence Closure

A family of local observation policies determines a collection of identity judgments distributed across a cover. Each local relation identifies states within one observational domain, and overlap compatibility ensures that no direct disagreement occurs where two domains intersect. The remaining problem is to determine which global equivalence relation is forced by those local judgments.

The most immediate construction is to combine all local relations into one relation on the total state space. If any observer identifies two states, the combined relation records that identification. This union contains exactly the identity declarations made locally and no others.

The union is generally not an equivalence relation.

Reflexivity and symmetry are inherited under mild assumptions, but transitivity can fail whenever two consecutive identifications are supplied by different domains. A state may be locally equivalent to a second state in one region, while the second is equivalent to a third state in another region, without any local observer directly identifying the first and third states.

Any global equivalence relation containing the local policies must nevertheless identify the first and third states. Transitivity is not optional. It is part of the logical structure of identity.

The canonical global construction is therefore obtained by taking the equivalence closure of the union of all local relations. This closure adds every identification required by reflexivity, symmetry, and transitivity while adding no unnecessary relation beyond those consequences.

The resulting equivalence relation is the minimal globalization.

It always exists.

It contains every local policy.

It is contained in every other global equivalence relation that contains those policies.

Its equivalence classes are the smallest globally coherent objects compatible with all declared local identities.

The construction solves the existence problem for a global relation containing the local data. It does not yet solve the faithful descent problem. The generated relation may introduce additional identities inside a local domain. When that occurs, globalization exists, but the original local policies are not recovered exactly.

This section develops the generated global relation in full. It begins with the union of local identities, shows why that union generally fails to be an equivalence relation, constructs its equivalence closure, characterizes the closure through finite zig-zags, proves minimality,

establishes a universal property, and interprets the resulting quotient as the global object space forced by the distributed observation system.

4.1 Union of Local Identity Relations

Let X be a state space.

Let $\{U_i\}_{i \in I}$ be a cover of X.

For each i, let E_i be an equivalence relation on U_i.

Each E_i is formally a subset of $U_i \times U_i$. Because $U_i \times U_i$ is contained in $X \times X$, every local relation may also be regarded as a relation on X whose domain of nontrivial declaration is U_i.

Define the union relation R on X by

$R = \bigcup_{i \in I} E_i$.

For $x, y \in X$,

$x \, R \, y$

exactly when there exists some $i \in I$ such that $x, y \in U_i$ and

$x \, E_i \, y$.

Thus R records every identity judgment made by at least one local observer.

The relation R is the raw global identity data.

It contains only explicit local identifications. It does not yet include identities implied by chains passing through several domains.

Because the domains cover X, every state belongs to at least one U_i. Since E_i is reflexive on U_i, every $x \in X$ satisfies

$x \, E_i \, x$

for some i.

Therefore

$x \, R \, x$.

The union relation is reflexive.

Each E_i is symmetric. If $x \, R \, y$, then $x \, E_i \, y$ for some i, and hence $y \, E_i \, x$. Therefore $y \, R \, x$.

The union relation is symmetric.

Transitivity is different.

Suppose x R y and y R z.

There exist indices i and j such that

x Ei y

and

y Ej z.

If i = j, then all three states lie in Ui, and transitivity of Ei gives

x Ei z.

Thus x R z.

If i and j are different, no single local relation need contain both identifications. The state y acts as a bridge between domains. There may be no observer whose domain contains both x and z, and even if such a domain exists, its policy may distinguish them.

Therefore x R z need not hold.

The union relation is generally a tolerance-like relation: reflexive and symmetric but not necessarily transitive.

This failure is not repaired by pairwise compatibility.

Compatibility ensures that Ei and Ej agree on Ui ∩ Uj. It says that if x and z both lie in the overlap, their identity judgment is consistent. It does not require an observer to identify endpoints connected through a chain whose intermediate states lie in different domains.

The union relation should therefore be interpreted as the graph of direct local identity.

Its vertices are states in X.

An undirected edge joins x and y whenever at least one local policy identifies them.

Reflexive loops occur at every vertex, although they may be omitted from the graphical representation.

Global equivalence will arise from connectivity in this graph rather than from adjacency alone.

This distinction is fundamental.

Direct identity is supplied by local observation.

Global identity is generated by compositional closure.

The relation R is also monotone in the local data. If every Ei is refined to a relation Fi with

$E_i \subseteq F_i$,

then the corresponding union grows:

$\bigcup_i E_i \subseteq \bigcup_i F_i$.

Adding local identifications can only add potential global identifications. Once transitive closure is taken, a small local change may create a much larger global fusion.

The union therefore contains the seeds of globalization, but its consequences may be highly nonlocal.

4.2 Failure of the Union to Be an Equivalence Relation

The simplest failure of transitivity occurs with three states.

Let

$X = \{a, b, c\}$.

Let

$U1 = \{a, b\}$

and

$U2 = \{b, c\}$.

Let E1 identify a and b.

Let E2 identify b and c.

Then

a R b

and

b R c.

But no local policy directly identifies a and c, because no domain contains both. Therefore

a is not R-related to c.

The relation R is not transitive.

Any global equivalence relation containing E1 and E2 must identify all three states. Thus the global object generated by the two local objects is larger than either local class.

This example represents legitimate identity propagation when no local policy objects to the result.

A more important example includes a return to an earlier domain.

Let

X = {a, b, c, d}.

Take

U1 = {a, b, d},

U2 = {b, c},

U3 = {c, d}.

Let E1 have classes

{a, b}

and

{d}.

Let E2 identify b and c.

Let E3 identify c and d.

The union contains

a R b,

b R c,

c R d.

It does not contain

a R d

because E1 distinguishes a and d.

Nevertheless, any global equivalence containing R must identify a and d through transitivity.

The union is therefore incomplete as a global identity relation, while its completion will conflict with E1.

This example separates two issues.

The failure of transitivity means R is not yet a global equivalence.

The forced identification of a and d means no faithful globalization is possible.

The first problem is solved by equivalence closure.

The second is revealed by restricting the closure back to U1.

The union may also fail transitivity through longer chains. A relation can be locally coherent at every step while globally requiring identifications across arbitrarily many domains.

For a finite cover, these chains are finite because every derivation of equivalence can be represented by a finite sequence of relation steps.

For an infinite state space or cover, the equivalence closure is still generated by finite chains. Equivalence relations are finitary in the sense that transitivity combines finitely many successive relations at a time. No genuinely infinite path is required to witness membership in the generated equivalence relation.

The failure of R to be transitive is therefore completely characterized by finite zig-zags.

The union R is itself an equivalence relation exactly when it is transitive.

This occurs when every composable pair of local identities already has a direct local identity consequence:

For all $x, y, z \in X$,

if $x \, E_i \, y$ and $y \, E_j \, z$ for some i and j,

then $x \, E_k \, z$ for some k.

When this condition holds, no new identities are added by closure:

$EqCl(R) = R$.

Such systems may be called transitively saturated local systems.

Transitive saturation is stronger than overlap compatibility. It requires the family of local policies, taken together, to be closed under all cross-domain compositions.

In many applications, this condition is too strong to expect at the level of direct observation. Global object formation exists precisely because identity can be inferred through chains not visible to any one observer.

The aim is therefore not to avoid closure but to control its effects.

4.3 Reflexive, Symmetric, and Transitive Closure

For an arbitrary relation S on X, its equivalence closure is the smallest equivalence relation containing S.

Write

EqCl(S)

for this closure.

Several equivalent constructions are available.

One construction defines EqCl(S) as the intersection of all equivalence relations on X that contain S:

$\mathrm{EqCl}(S) = \bigcap\{E \in \mathrm{Eq}(X) : S \subseteq E\}$.

This family is nonempty because the universal relation $X \times X$ is an equivalence relation containing every relation S.

An arbitrary intersection of equivalence relations is again an equivalence relation. Therefore EqCl(S) is an equivalence relation.

It contains S because every relation in the intersection contains S.

It is minimal because any equivalence relation containing S appears among the relations being intersected, and the intersection is contained in each member of the family.

This definition proves existence and minimality immediately.

A second construction proceeds operationally.

Begin with S.

Add every pair (x, x) to enforce reflexivity.

Add the reverse of every pair to enforce symmetry.

Then repeatedly add (x, z) whenever (x, y) and (y, z) are present, enforcing transitivity.

Continue until no new pairs are required.

For finite X, this process terminates after finitely many additions because $X \times X$ is finite.

For infinite X, the closure is still well defined, although a naive iterative procedure may require transfinite bookkeeping if formulated as successive relation stages. The finite-chain characterization avoids this complication.

In the present setting, the union R is already reflexive and symmetric because the domains cover X and every Ei is an equivalence relation. Therefore only transitive closure is needed:

$E^* = EqCl(R) = TransCl(R)$.

The notation E^* will denote the generated global equivalence relation.

For $x, y \in X$,

$x \; E^* \; y$

exactly when x and y are connected by a finite R-chain.

Because R is symmetric, the chain may move through local identities in either direction. Since each Ei is already symmetric, no separate orientation is needed.

The closure operator EqCl has several standard properties.

It is extensive:

$S \subseteq EqCl(S)$.

It is monotone:

$S \subseteq T$ implies $EqCl(S) \subseteq EqCl(T)$.

It is idempotent:

$EqCl(EqCl(S)) = EqCl(S)$.

These properties make equivalence closure a closure operator on the lattice of binary relations on X.

Extensivity states that globalization preserves every declared local identity.

Monotonicity states that adding local identities cannot remove global identities.

Idempotence states that once global equivalence has been formed, repeating the globalization adds nothing.

The fixed points of EqCl are precisely the equivalence relations.

Thus equivalence relations are the closed identity systems, while arbitrary reflexive symmetric relations are partially specified systems awaiting transitive completion.

The generated globalization is therefore the closure of distributed local identity under the laws required for coherent objecthood.

4.4 Minimal Equivalence Generated by Local Policies

Apply equivalence closure to the union of local policies:

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$.

This relation will be called the generated global equivalence, the minimal globalization, or the equivalence generated by the local system.

It satisfies three defining properties.

First,

$E_i \subseteq E^*$

for every i.

Every locally licensed identity is globally preserved.

Second, E^* is an equivalence relation on X.

It therefore defines a legitimate global quotient.

Third, if F is any equivalence relation on X such that

$E_i \subseteq F$

for every i,

then

$E^* \subseteq F$.

Thus E^* adds no identification that is not forced by equivalence closure of the local declarations.

These properties characterize E^* uniquely.

Suppose E' is another equivalence relation satisfying the same three conditions. Since E' contains every E_i, minimality of E^* gives

$E^* \subseteq E'$.

Since E^* also contains every E_i and E' is assumed minimal, one obtains

$E' \subseteq E^*$.

Therefore

$E' = E^*$.

The generated relation is canonical.

It requires no arbitrary choice of representatives, ordering of domains, or sequence of pairwise mergers. Any correct globalization procedure that produces the smallest containing equivalence must yield E*.

The relation E* can also be expressed as the join of the local relations after they are embedded into the lattice Eq(X).

A technical point is required because each Ei is originally defined only on Ui. To regard it as a relation on all of X, one may extend it by equality outside Ui:

$\hat{E}_i = E_i \cup \Delta_{X \setminus U_i}$,

together with the identity pairs already present on Ui.

More simply, because the union over a cover already contains every diagonal pair, one may work directly with the subsets $E_i \subseteq X \times X$.

Then

$E^* = \bigvee_i E_i$

in the sense of the equivalence relation generated by their union.

This is a global join across domains.

The closure may join states that no observer sees together. That is not an accidental overreach. It is the precise meaning of global equivalence generated by distributed identity.

The minimality of E* does not mean that its equivalence classes are small in cardinality. A single bridging relation can cause several large classes to merge. Minimality concerns inclusion among relations, not the number or size of quotient classes.

A minimal addition of one pair may have extensive transitive consequences.

This sensitivity is central to distributed identity systems. Entity-resolution networks, genealogical classifications, sensor associations, semantic graphs, and scientific integrations can undergo abrupt global fusion when one new bridge closes a path.

The generated relation provides the exact mathematical representation of that effect.

4.5 Zig-Zag Paths of Identification

The most useful concrete characterization of E* is through finite zig-zags.

A local identity zig-zag from x to y is a finite sequence

$x = x_0, x_1, \ldots, x_n = y$

such that for every k with $1 \le k \le n$, there exists an index ik satisfying

$x_{k-1} \; E_{i_k} \; x_k$.

The domains may change at every step.

The intermediate state xk must belong to both domains used in the adjacent steps whenever those steps use different local policies. It acts as a bridge through the observational cover.

The zig-zag theorem states:

For $x, y \in X$,

$x \; E^* \; y$

if and only if there exists a finite local identity zig-zag from x to y.

To prove the forward direction constructively, define a relation Z by

$x \; Z \; y$

when a finite local identity zig-zag connects x to y.

The relation Z contains every Ei because a single local equivalence gives a zig-zag of length one.

The relation Z is reflexive because the sequence of length zero connects x to itself.

It is symmetric because a zig-zag can be reversed, using symmetry of every Ei.

It is transitive because two zig-zags can be concatenated.

Therefore Z is an equivalence relation containing all Ei.

By minimality of E*,

$E^* \subseteq Z$.

Conversely, because E* contains every Ei and is transitive, every finite zig-zag has E*-equivalent endpoints. Thus

$Z \subseteq E^*$.

Therefore

$Z = E^*$.

This characterization converts global equivalence into a path problem.

Construct an undirected graph G whose vertex set is X. Join x and y by an edge whenever $x \; E_i \; y$ for at least one i. Then the E*-equivalence classes are exactly the connected components of G.

For finite systems, minimal globalization can therefore be computed using standard connected-component algorithms or disjoint-set union structures.

The graph may be compressed further by using local quotient classes as nodes. Each Ei-class becomes a local-object node. Two nodes from different domains are connected when they share at least one state. Connected components in this local-object graph correspond to E*-classes.

The zig-zag view also yields explicit obstruction witnesses.

Suppose x and y both lie in Ui, and Ei distinguishes them. If a zig-zag through the cover connects x and y, then

x E* y

while

x is not Ei-equivalent to y.

The zig-zag is a certificate that faithful descent fails on Ui.

A shortest such zig-zag is a minimal obstruction witness.

Its length measures how many local identity steps are required to produce the forbidden fusion.

The number of distinct domains traversed measures the distributed depth of the obstruction.

A short obstruction may indicate an easily identifiable bridge error.

A long obstruction may reveal a genuinely emergent global relation that becomes visible only through extensive integration.

The mathematics alone records the path and the resulting identity. Interpretation determines whether the path represents valid evidence, ambiguous association, or destructive propagation.

4.6 The Minimal Globalization Theorem

The basic construction can now be stated formally.

Minimal Globalization Theorem.

Let X be a set, let $\{U_i\}_{i \in I}$ be a cover of X, and let E_i be an equivalence relation on U_i for each i. Define

$R = \bigcup_{i \in I} E_i$

and

$E^* = \mathrm{EqCl}(R)$.

Then:

E^* is an equivalence relation on X.

For every i,

$E_i \subseteq E^* | U_i$.

If F is any equivalence relation on X satisfying

$E_i \subseteq F | U_i$

for every i, then

$E^* \subseteq F$.

For $x, y \in X$,

$x \; E^* \; y$

if and only if there is a finite sequence

$x = x_0, x_1, \ldots, x_n = y$

such that each consecutive pair belongs to some local relation E_i.

Proof.

Because $\mathrm{EqCl}(R)$ is defined as the intersection of all equivalence relations containing R, E^* is an equivalence relation on X and contains R.

Since $E_i \subseteq R$ for every i, one has

$E_i \subseteq E^*$.

Restricting to U_i gives

$E_i \subseteq E^* | U_i$.

Let F be an equivalence relation on X containing every local relation. Then $R \subseteq F$. By minimality of equivalence closure,

$\mathrm{EqCl}(R) \subseteq F$.

Therefore

$E^* \subseteq F$.

The zig-zag characterization follows from the construction in the previous subsection. The relation of finite zig-zag connectivity is an equivalence relation containing R, and every equivalence relation containing R must contain the endpoints of every finite R-chain. Hence it equals E*.

This proves the theorem.

The theorem makes no compatibility assumption.

Even incompatible local policies generate a minimal global equivalence relation containing all local identifications. If two observers disagree on an overlap, the generated relation accepts the coarser judgment: any pair identified by either observer becomes globally identified.

Thus minimal globalization always resolves conflict by retaining all positive identity declarations. It cannot preserve a local declaration of distinction against another observer's declaration of equivalence because distinction is represented by absence from a relation, while equivalence is represented positively by inclusion.

This asymmetry is important.

Equivalence closure treats identity declarations as constraints that must be satisfied.

Local distinctions are fidelity conditions tested afterward.

The generated relation is therefore permissive toward identity and potentially destructive toward distinction.

A different conflict-resolution policy could reject some local identities before globalization. Such policy selection lies beyond the minimal construction and will become part of a later admissibility theory.

The theorem answers only:

What is the smallest coherent global identity relation accepting every local identification?

Its answer is E*.

4.7 Universal Property of Generated Globalization

The minimal globalization possesses a universal property that characterizes it independently of its explicit construction.

Let $q^*: X \to X/E^*$ be the canonical quotient map.

For every i, the local quotient map

$q_i: U_i \to U_i/E_i$

identifies states according to Ei.

Because Ei ⊆ E*, the restriction of q* to Ui is constant on every Ei-class. Therefore it factors uniquely through qi.

There exists a unique map

$\phi i: Ui/Ei \to X/E^*$

such that

$q^*|Ui = \phi i \circ qi$.

Explicitly,

$\phi i([x]Ei) = [x]E^*$.

This map is well defined because x Ei y implies x E* y.

The maps ϕi send local objects into their generated global objects.

Several distinct local objects may map to the same global object. This occurs when a zig-zag connects them.

If two distinct Ei-classes map to the same E*-class, then globalization has fused objects within the same local domain. This is precisely a descent defect.

Now let F be any equivalence relation on X containing every Ei, with quotient map

$qF: X \to X/F$.

Because E* ⊆ F, there is a unique canonical surjection

$\psi: X/E^* \to X/F$

such that

$qF = \psi \circ q^*$.

It is defined by

$\psi([x]E^*) = [x]F$.

Thus every other quotient accepting all local identities factors uniquely through the minimal global quotient.

The diagram is

$X \to X/E^* \to X/F$.

The quotient X/E* is therefore the finest global quotient compatible with every local identity declaration.

It preserves the maximum amount of distinction possible while still accepting all local equivalences.

Any coarser global quotient loses at least as much information.

This is the universal property of generated globalization:

X/E* is initial among global quotient object spaces whose quotient maps identify every locally equivalent pair.

Equivalently, E* is the least global identity policy containing all local policies.

At the level of maps, suppose $h: X \to Y$ is any function satisfying

$x \; E_i \; y$ implies $h(x) = h(y)$

for every i.

Then h is constant along every local identity step. It is therefore constant along every finite zig-zag. Hence

$x \; E^* \; y$ implies $h(x) = h(y)$.

There exists a unique function

$\bar{h}: X/E^* \to Y$

such that

$h = \bar{h} \circ q^*$.

This is a familiar quotient universal property, but its distributed interpretation is significant.

Any representation, classifier, invariant, or transformation that treats every locally equivalent pair identically must also treat every globally generated pair identically.

A system cannot accept all local substitutions while rejecting one of their transitive consequences.

The generated quotient is therefore the unavoidable semantic target of all globally consistent maps respecting the local policies.

This universal property gives the minimal globalization an operational meaning beyond set-theoretic closure. It is the object space through which every locally identity-preserving representation must factor.

4.8 Global Quotient Objects Produced by Closure

The quotient

X/E^*

is the global object space generated by the distributed observation system.

Its elements are the E^*-equivalence classes.

Each global class is a connected component of the local identity network.

A global object may have several forms.

It may be entirely contained in one domain and coincide with a local object.

It may extend across several domains through shared states.

It may combine local objects from different observers into one distributed object.

It may fuse several distinct objects belonging to one local observer after a zig-zag leaves and returns to that observer's domain.

The first three possibilities can represent successful object integration.

The fourth signals a failure of faithful descent unless the local policy is intentionally revised.

A global class $C \subseteq X$ has a local trace on U_i:

$C \cap U_i$.

If descent is faithful, every nonempty trace is exactly one E_i-class.

If descent is not faithful, a trace may be a union of several E_i-classes.

Thus a generated global object can appear locally as one object or as several distinct objects depending on whether closure respects the local resolution.

The map

$\phi_i: U_i/E_i \to X/E^*$

assigns every local object to its global component.

Faithful descent requires ϕ_i to be injective for every i.

Surjectivity is not required. A global object may fail to intersect a given domain, so no object in U_i/E_i maps to it.

Therefore faithful descent at the quotient level is equivalent to every local-to-global object map being injective.

This formulation is useful.

Local classes may merge across different domains; that is the purpose of gluing.

They may not merge within the same domain if the original local policy is to be preserved.

A generated global object may also be emergent. No single observer may possess a class containing all its states. The object exists only as the closure of a distributed chain.

For example, a global trajectory may be assembled from local track segments.

A biological population may be assembled from overlapping reproductive communities.

A scientific entity may be assembled from partial records distributed across databases.

A semantic concept may be assembled from overlapping contextual uses.

A physical object may be reconstructed from local coordinate descriptions.

In each case, the global object is not a mere union chosen arbitrarily. It is the equivalence class forced by the network of local identities.

The distinction between generated and arbitrary union is crucial.

A global object consists only of states connected through licensed local equivalences.

Every pair of states in the class is related by some finite chain of identity judgments.

The object therefore has an evidential or structural path supporting its unity.

However, the presence of a path does not prove that the resulting object is appropriate for every task. A mistaken local edge can merge entire components. A valid identity under one purpose may be destructive under another. The global quotient is always relative to the supplied policies.

Generated globalization should therefore be understood as a logical completion, not an infallible discovery mechanism.

It answers what follows from the local identity declarations under equivalence.

It does not decide whether those declarations were scientifically, computationally, or ontologically justified.

The construction nevertheless establishes a rigorous baseline.

Every distributed object system accepting all local identities must be at least as coarse as X/E*.

Every faithful globalization, if one exists, must contain the same generated equivalence.

Every descent obstruction is already visible in the restrictions of E*.

The local-to-global problem can therefore be reduced to a canonical test:

Construct the minimal globalization.

Restrict it back to every observational domain.

Compare those restrictions with the original local policies.

If every restriction agrees exactly, the local objects descend faithfully.

If any restriction is coarser, transitive globalization has created a forbidden local fusion.

The next section states and proves this faithful descent criterion. It shows that no separate search over all possible global equivalence relations is necessary. The minimal globalization alone determines whether exact local recovery is possible.

5. The Faithful Descent Theorem

The minimal globalization always produces a global equivalence relation containing every local identity judgment. It therefore guarantees that distributed local policies can be embedded into some coherent global identity system. What it does not guarantee is that the original local distinctions survive.

Faithful descent is the stronger requirement.

A global policy is faithful to the local system when restricting it to each observational domain reproduces the local policy exactly. It preserves every local identification and every local distinction.

The distinction between containment and equality is decisive.

For every local policy Ei, minimal globalization automatically gives

$E_i \subseteq E^*|U_i$.

Thus every locally equivalent pair remains globally equivalent.

Faithfulness additionally requires the reverse inclusion:

$E^*|U_i \subseteq E_i$.

No pair distinguished locally may become identified through a global zig-zag.

The main theorem of this paper states that faithful descent can be tested entirely through the minimal globalization. There is no need to search through all possible global equivalence relations. If the generated global relation introduces a forbidden local identity, every global relation containing the local policies must introduce it as well. If the generated relation introduces no forbidden identity, it is itself a faithful globalization.

This produces a complete existence criterion.

The theorem has several equivalent forms.

It can be stated through restriction of the generated relation.

It can be stated through zig-zag closure.

It can be stated through injectivity of local-to-global quotient maps.

It can be stated through the condition that no global connected component contains two distinct local objects from the same domain.

Each form emphasizes a different aspect of distributed object formation.

The restriction form is best suited to order theory and categorical formulation.

The zig-zag form exposes explicit obstruction witnesses.

The quotient form clarifies object preservation.

The connected-component form supports finite computation.

Together they provide the mathematical core of faithful local-to-global objecthood.

5.1 Exact Recovery of Local Policies

Let X be a state space covered by domains $\{U_i\}_{i \in I}$.

Let E_i be an equivalence relation on U_i.

A global equivalence relation E on X is a faithful globalization of the local family when

$E|U_i = E_i$

for every i.

This equality contains two conditions.

First,

$E_i \subseteq E|U_i$.

Every local identity is accepted globally.

Second,

$E|U_i \subseteq E_i$.

Every global identity between states jointly visible in Ui was already accepted locally.

The first condition is identity preservation.

The second condition is distinction preservation.

A global relation satisfying only the first condition may be called an extending globalization. It extends every local policy but may collapse additional distinctions.

The minimal globalization E* is always an extending globalization.

Faithful globalization requires exact recovery.

Exact recovery can be interpreted operationally. Suppose an observer first forms its local quotient Ui/Ei. The distributed system then constructs a global quotient X/E. If the global quotient is restricted back to the states in Ui, the observer should recover precisely its original partition.

No two local objects should have fused.

No local object should have split, although splitting is already impossible when E contains Ei.

Because global equivalence extends local equivalence, the only possible local distortion is fusion.

Thus faithful descent is fundamentally a no-unlicensed-fusion condition.

At the quotient level, define the local-to-global map

$\phi_{iE}: U_i/E_i \to X/E$

by

$\phi_{iE}([x]_{E_i}) = [x]_E$.

This map is well defined exactly when $E_i \subseteq E|U_i$.

The map ϕ_{iE} is injective exactly when

$E|U_i = E_i$.

To see this, suppose ϕiE is injective. If x E y with x, y ∈ Ui, then

ϕiE([x]Ei) = [x]E = [y]E = ϕiE([y]Ei).

Injectivity gives

[x]Ei = [y]Ei,

so x Ei y.

Thus E|Ui ⊆ Ei.

The reverse inclusion already follows from well-definedness.

Conversely, if E|Ui = Ei and

ϕiE([x]Ei) = ϕiE([y]Ei),

then x E y. Because x and y lie in Ui, exact recovery gives x Ei y, so the two local classes are equal.

Therefore faithful descent is equivalent to the injectivity of every local-to-global quotient map.

This formulation gives a clear object-level meaning:

Every local object remains distinct from every other local object belonging to the same observational domain.

Objects from different domains may still merge globally. That merging is the purpose of gluing.

Faithfulness forbids only those mergers that contradict an existing local distinction.

5.2 Restricting the Generated Global Relation

Let

R = ⋃i Ei

and

E* = EqCl(R).

Because Ei ⊆ E*, one always has

Ei ⊆ E*|Ui.

The restriction E*|Ui may therefore equal Ei or be strictly coarser.

If

E*|Ui = Ei,

then the generated global identity introduces no new relation inside Ui.

If

$E_i \subset E^*|U_i$,

then there exist $x, y \in U_i$ such that

$x \; E^* \; y$

but

x is not Ei-equivalent to y.

The states x and y are locally distinct but globally connected through a finite sequence of local identities.

This strict inclusion is the precise signature of descent failure on Ui.

The restriction test is canonical because E* is contained in every global equivalence relation that extends all local policies.

Suppose F is any global equivalence relation with

$E_i \subseteq F|U_i$

for all i.

Then

$E^* \subseteq F$.

Therefore

$E^*|U_i \subseteq F|U_i$.

Any local fusion already produced by E* must also occur under F.

Consequently, once E* identifies a locally distinguished pair, no larger global equivalence can repair the damage. Coarsening can only introduce additional identifications.

Faithful descent can fail only in one direction. A candidate global relation cannot be made more faithful by adding equivalences. It can be repaired only by changing the local data, rejecting some local identity declarations, refining the cover, or altering the meaning of admissible globalization.

This monotonicity makes E* the decisive test object.

If E* is faithful, it solves the descent problem.

If E* is unfaithful, no extending globalization can solve it.

The test can also be written as a saturation condition.

A subset Ui is saturated relative to E* at the level of local classes when every intersection of an E*-class with Ui is exactly one Ei-class or is empty.

Equivalently, for every x ∈ Ui,

[x]E* ∩ Ui = [x]Ei.

This equality says that the global object containing x has precisely the locally licensed trace inside Ui.

It may extend beyond Ui, but it may not return to Ui through another locally distinct class.

This return condition is the central geometric image of faithful descent.

A global object may pass through many domains.

It may intersect a given domain only in one local object.

If it intersects the same domain in several locally distinct objects, descent fails.

5.3 Statement of the Faithful Descent Criterion

Faithful Descent Theorem.

Let X be a set, let {Ui}i∈I be a cover of X, and let Ei be an equivalence relation on Ui for every i. Define

E* = EqCl(⋃i∈I Ei).

Then the following conditions are equivalent.

First, there exists an equivalence relation E on X such that

E|Ui = Ei

for every i.

Second,

E*|Ui = Ei

for every i.

Third, for every i and every x, y ∈ Ui, if there exists a finite local identity zig-zag

x = x0, x1, …, xn = y

with each consecutive pair related by some Ej, then

x Ei y.

Fourth, for every i, the canonical map

$\phi i: Ui/Ei \to X/E^*$

defined by

$\phi i([x]Ei) = [x]E^*$

is injective.

Fifth, no E*-equivalence class intersects any Ui in more than one Ei-equivalence class.

Sixth, in the graph of local quotient objects and overlap correspondences, no connected component contains two distinct nodes belonging to the same local quotient Ui/Ei.

When these conditions hold, E* is a faithful globalization. It is the unique smallest faithful globalization under inclusion of equivalence relations.

The theorem requires no separate assumption of pairwise overlap compatibility. Compatibility follows automatically from any of the equivalent conditions.

Indeed, if E* restricts to both Ei and Ej, then on $Ui \cap Uj$,

$Ei|Ui\cap Uj = E^*|Ui\cap Uj = Ej|Ui\cap Uj$.

Thus faithful descent implies overlap agreement.

The theorem also clarifies the role of compatibility. Pairwise agreement is an efficiently testable necessary condition, but the full criterion is closure under every returning zig-zag.

The local family descends exactly when all global identity paths are already locally valid wherever their endpoints become jointly observable.

This is the algebraic law of faithful distributed object formation.

5.4 Necessity of Local Recovery

Assume there exists a faithful globalization E on X satisfying

$E|Ui = Ei$

for every i.

Because E contains each Ei, it contains their union:

$\bigcup_i E_i \subseteq E$.

Since E is an equivalence relation, it contains the equivalence closure of that union:

$E^* \subseteq E$.

Restricting to Ui gives

$E^*|U_i \subseteq E|U_i$.

By faithfulness,

$E|U_i = E_i$.

Therefore

$E^*|U_i \subseteq E_i$.

But minimal globalization always contains the local relation:

$E_i \subseteq E^*|U_i$.

Combining the two inclusions gives

$E^*|U_i = E_i$.

Thus exact recovery by E* is necessary.

The argument is short, but its interpretation is important.

Every identity generated by local zig-zags is logically unavoidable in any global equivalence accepting the local policies.

A faithful global relation cannot choose to ignore one of those consequences.

If a zig-zag connects two states in Ui, the global relation E must identify them.

Because E restricts exactly to Ei, the local observer must already identify them.

Therefore every returning global identity path must close locally.

The necessity proof can be expressed directly in zig-zag form.

Let $x, y \in U_i$ and suppose

$x = x_0, x_1, \ldots, x_n = y$

is a local identity zig-zag.

Every consecutive relation belongs to some Ej and hence to E.

By transitivity of E,

x E y.

Because x and y lie in Ui and E|Ui = Ei,

x Ei y.

Thus the zig-zag criterion is necessary.

At the quotient level, assume E is faithful. The map

ϕiE: Ui/Ei → X/E

is injective.

Because E* ⊆ E, the canonical map X/E* → X/E factors the local object maps. If two distinct Ei-classes merged already in X/E*, they would also merge in X/E, contradicting injectivity.

Therefore the maps Ui/Ei → X/E* must also be injective.

All formulations express the same fact:

Minimal globalization contains only identities that every possible globalization must contain.

Its local distortions are therefore universal obstructions.

5.5 Sufficiency of the Restriction Condition

Assume

E*|Ui = Ei

for every i.

The relation E* is already an equivalence relation on X.

By assumption, its restriction to every local domain is exactly the corresponding local policy.

Therefore E* is a faithful globalization.

This proves sufficiency.

The simplicity of this direction is one of the strengths of the construction. Once the minimal globalization has been formed, no additional gluing choices are required. Exact local recovery immediately certifies existence.

The zig-zag form follows directly.

Assume that every local identity zig-zag whose endpoints lie together in Ui has Ei-equivalent endpoints.

Take x, y ∈ Ui with x E* y.

By the zig-zag characterization of E*, there exists a finite local identity path from x to y.

The assumed condition gives x Ei y.

Therefore

E*|Ui ⊆ Ei.

Since Ei ⊆ E*|Ui always holds,

E*|Ui = Ei.

Thus E* is faithful.

The quotient-map form is also sufficient.

Assume every canonical map

ϕi: Ui/Ei → X/E*

is injective.

Take x, y ∈ Ui with x E* y.

Then

ϕi([x]Ei) = [x]E* = [y]E* = ϕi([y]Ei).

Injectivity gives

[x]Ei = [y]Ei.

Hence x Ei y.

Again,

E*|Ui = Ei.

The connected-component form follows because E*-classes are the connected components of the local identity graph. If no component contains two distinct Ei-classes, then each component intersects Ui in at most one local class. Since every local class lies in some component, the intersection is exactly that class whenever nonempty.

Each criterion therefore supplies a complete algorithmic test.

Construct the local identity graph or the local-class correspondence graph.

Compute its connected components.

For each domain Ui, check whether two distinct Ei-classes occur in the same component.

If none do, descent succeeds.

If some do, their component supplies a witness of failure.

5.6 Proof of the Faithful Descent Theorem

Proof.

Let

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$.

The equivalence of the six conditions will be established in sequence.

Assume the first condition: there exists an equivalence relation E on X with

$E|U_i = E_i$

for every i.

Because E contains every Ei, minimality of E* gives

$E^* \subseteq E$.

Therefore

$E^*|U_i \subseteq E|U_i = E_i$.

Since $E_i \subseteq E^*|U_i$, equality follows:

$E^*|U_i = E_i$.

Thus the first condition implies the second.

Assume the second condition. Because E* is an equivalence relation on X and restricts to Ei on every Ui, it is itself a faithful globalization. Thus the second condition implies the first.

The first two conditions are equivalent.

Assume the second condition. Let $x, y \in U_i$ and suppose a finite local identity zig-zag connects x to y. By the zig-zag characterization,

$x \; E^* \; y$.

Since $E^*|U_i = E_i$,

x Ei y.

Thus the second condition implies the third.

Assume the third condition. Let x, y ∈ Ui satisfy x E* y. A finite local identity zig-zag connects them. By the third condition,

x Ei y.

Hence

E*|Ui ⊆ Ei.

The reverse inclusion is automatic, so

E*|Ui = Ei.

Thus the third condition implies the second.

The second and third conditions are equivalent.

Assume the second condition. Suppose

ϕi([x]Ei) = ϕi([y]Ei).

Then

[x]E* = [y]E*,

so x E* y.

Since x and y lie in Ui and E*|Ui = Ei,

x Ei y.

Therefore

[x]Ei = [y]Ei.

Thus ϕi is injective, and the second condition implies the fourth.

Assume the fourth condition. If x, y ∈ Ui and x E* y, then

ϕi([x]Ei) = [x]E* = [y]E* = ϕi([y]Ei).

Injectivity gives

[x]Ei = [y]Ei,

so x Ei y.

Therefore $E^*|U_i \subseteq E_i$. Again equality follows.

Thus the fourth condition implies the second.

The second and fourth conditions are equivalent.

Assume the second condition. Let C be an E^*-class. If $C \cap U_i$ is nonempty, choose $x \in C \cap U_i$. Then

$C \cap U_i = [x]E^* \cap U_i$.

Because $E^*|U_i = E_i$,

$[x]E^* \cap U_i = [x]E_i$.

Thus C intersects U_i in exactly one E_i-class.

Therefore the second condition implies the fifth.

Assume the fifth condition. Let $x, y \in U_i$ satisfy $x\ E^*\ y$. They belong to the same E^*-class C. Both lie in $C \cap U_i$. By the fifth condition, $C \cap U_i$ is contained in one E_i-class. Hence $x\ E_i\ y$.

Thus $E^*|U_i \subseteq E_i$, and equality follows.

The fifth condition implies the second.

The second and fifth conditions are equivalent.

Finally, the E^*-classes are precisely the connected components of the local-object correspondence graph after states within each E_i-class are collapsed to one node. A component contains two distinct nodes from U_i/E_i exactly when one E^*-class intersects U_i in at least two E_i-classes.

Therefore the fifth and sixth conditions are equivalent.

All six conditions are equivalent.

When they hold, E^* is faithful. If F is any faithful globalization, then F contains every E_i and therefore contains E^*. Hence E^* is contained in every faithful globalization. It is the unique smallest faithful globalization.

This proves the theorem.

5.7 Uniqueness and Minimality of Faithful Globalization

The faithful descent theorem establishes that E^* is the unique smallest faithful globalization.

It does not imply that every faithful globalization equals E^*.

This distinction is essential.

A global equivalence relation F may strictly contain E* while still restricting to Ei on every Ui. Such an F may identify states that never occur together in any local domain whose policy distinguishes them.

For example, suppose X is covered by two disjoint domains U1 and U2. The local policies impose no relation between a state in U1 and a state in U2.

The minimal globalization keeps the two components separate.

A coarser global relation may identify one object from U1 with one object from U2. Because no local domain contains both objects, this new identity does not alter either local restriction.

Both global relations are faithful to the local data.

The local system underdetermines cross-component identity.

Thus faithful globalizations may form a family.

E* is the finest member of that family. It makes only the identities forced by local data.

Other faithful globalizations add identities not decided locally.

The set of faithful globalizations can be written as

$\mathrm{Faith}(\{E_i\}) = \{F \in \mathrm{Eq}(X) : F|U_i = E_i \text{ for all } i\}$.

When nonempty, this set has E* as its least element under refinement.

It may or may not have a greatest element.

Arbitrary joins of faithful globalizations need not remain faithful, because combining separately harmless cross-domain identities can create a zig-zag that returns to a domain and fuses local objects.

Therefore the family of faithful globalizations is not generally closed under joins.

It is closed under arbitrary intersections.

If $\{F_\alpha\}$ is a family of faithful globalizations, then

$F = \bigcap_\alpha F_\alpha$

is an equivalence relation.

On Ui,

$F|U_i = \bigcap_\alpha (F_\alpha|U_i) = \bigcap_\alpha E_i = E_i$.

Thus F is faithful.

This closure under meets confirms again that E*, when faithful, is the intersection of all faithful globalizations and hence their least element.

Uniqueness of faithful globalization requires additional conditions.

One sufficient condition is pair coverage:

For every x, y ∈ X, there exists some Ui containing both x and y.

Under this condition, every possible pairwise identity is locally tested. If E and F are faithful globalizations, then for any x and y choose Ui containing both. One has

x E y exactly when x Ei y exactly when x F y.

Therefore E = F.

Pair coverage is strong. It effectively gives the local system complete pairwise observational access even if no single domain contains all states.

A weaker uniqueness condition can be stated at the object level.

If every pair of distinct E*-classes has representatives that occur together in some Ui, then no faithful globalization can merge those classes. Hence E* is the only faithful globalization.

In general, the space between minimality and uniqueness records unresolved global identity.

The local data may determine some global objects completely while leaving other cross-domain mergers undecided.

A mature distributed observation theory should distinguish:

forced identity,

forbidden identity,

and underdetermined identity.

Forced identity is represented by E*.

Forbidden identity is any relation between states jointly contained in a domain where the local policy distinguishes them.

Underdetermined identity concerns pairs neither forced together by local zig-zags nor forbidden by any local comparison.

Faithful globalizations differ only in how they resolve underdetermined identities, subject to the requirement that their transitive consequences remain locally faithful.

This trichotomy will be important in later theories of admissibility and policy selection.

5.8 Interpretation as Lawful Global Object Formation

The faithful descent theorem provides a precise criterion for when local objects can be interpreted as consistent views of global objects.

A distributed observation system forms lawful global objects when every global identity generated by local equivalences preserves all distinctions that were certified locally.

The theorem does not require every global object to appear completely in one domain.

A global object may be distributed across many observers.

It may be constructed only through a long chain of overlaps.

It may possess no privileged representative and no central observer.

Faithful descent requires only that the object never return to one domain in two forms that the local policy distinguishes.

This is the essential law:

A global object may extend local identity, but it may not contradict local identity.

The theorem also separates construction from evaluation.

Equivalence closure constructs the global objects forced by the local policies.

Restriction evaluates whether those objects are faithful to the policies that generated them.

This separation prevents a common conceptual error. One might assume that because a global equivalence relation can always be formed, local data always glue successfully. The theorem shows otherwise.

Globalization is automatic.

Faithful globalization is conditional.

The condition is not merely pairwise agreement. It is closure under every distributed identity path.

The result has direct implications across applications.

In entity resolution, local record matches can be merged safely exactly when no connected match component contains two records that one source explicitly distinguishes as separate entities.

In multisensor tracking, local tracks form global trajectories faithfully exactly when no global association loop merges two tracks distinguished by the same sensor.

In biological classification, regional classes descend exactly when chains of local equivalence do not force the merger of taxa distinguished within one region.

In scientific integration, concepts from different disciplines can be unified faithfully exactly when the resulting equivalence closure preserves every distinction maintained within each contributing framework.

In artificial intelligence, representations from specialized models can be combined faithfully exactly when cross-model associations do not collapse categories that any model must retain for its task.

In multiscale physics, local or scale-dependent objects descend exactly when the global identification structure does not erase distinctions required within any one effective theory.

The theorem therefore functions as both a mathematical result and a diagnostic principle.

It determines when distributed evidence supports one coherent object ontology.

It identifies the precise pairs responsible when coherence fails.

It reveals which global identities are unavoidable.

It separates locally justified fusion from globally induced overidentification.

The central sequence can now be stated in theorem form:

Local observers declare equivalences.

The union records direct identity.

Equivalence closure generates unavoidable global identity.

Restriction tests local fidelity.

Faithful descent occurs exactly when closure creates no new relation inside any local domain.

When the test succeeds, X/E* is the finest global object space realizing all local policies exactly.

When it fails, no global equivalence relation can accept every local identity while preserving every local distinction.

The next section develops the failure case. It defines gluing obstructions and descent defects, constructs minimal zig-zag witnesses, introduces finite measures of defect, and examines how an obstructed local system may be repaired.

6. Gluing Obstructions and Descent Defects

Faithful descent fails when globally generated identity becomes coarser than at least one local identity policy. The local relations may be individually valid and mutually compatible on every overlap, yet their transitive closure can still force two locally distinct states into one global equivalence class.

This failure is not merely an absence of a desired construction. It possesses explicit mathematical structure.

The generated global equivalence

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$

contains every identity forced by the local policies. For each domain U_i, the original local relation E_i is contained in the restriction $E^*|U_i$. The difference between them records exactly which additional identities have entered U_i through global closure.

Define the descent defect on U_i by

$D_i = E^*|U_i \setminus E_i$.

A pair (x, y) belongs to D_i when x and y are jointly visible in U_i, distinguished by the local observer, but connected globally by a finite chain of locally licensed identities.

The relation D_i is therefore the set of forbidden local fusions produced by globalization.

Faithful descent occurs exactly when every D_i is empty.

When some D_i is nonempty, each of its pairs has a zig-zag witness. The witness begins at one local object, leaves or traverses the observational network, and returns to the same domain through a different local object. This returning path forces a merger that the local policy does not authorize.

The defect relation identifies the endpoints of the failure.

The zig-zag identifies its mechanism.

The connected global component identifies the scale of the resulting fusion.

These three levels form the basis of an obstruction theory for distributed object formation.

A gluing obstruction should not be interpreted automatically as evidence that one observer is wrong. It may indicate an erroneous bridge, incompatible purposes, insufficiently expressive local policies, a cover that hides relevant comparisons, or a genuine global identity that requires revision of a local ontology. The mathematics isolates the conflict but does not prejudge its scientific interpretation.

This section defines obstruction relations and witnesses, develops quantitative measures for finite systems, distinguishes local compatibility from global fidelity, and describes several methods for repairing an obstructed observation network.

6.1 Identification Chains Crossing Observational Domains

Let x and y be states in the same domain Ui.

Suppose the local policy distinguishes them:

x is not Ei-equivalent to y.

If x and y are nevertheless related by E*, then there exists a finite sequence

x = x0, x1, …, xn = y

such that every consecutive pair is identified by some local policy:

xk−1 Eik xk.

The sequence may begin with an identification inside Ui, pass through several other domains, and eventually return to Ui. It may also leave Ui immediately through an overlap state and return only at the final step.

Because the endpoints lie in Ui but belong to different Ei-classes, the chain creates a returning identification.

A returning identification path is therefore a local identity zig-zag whose endpoints lie in one domain but are locally inequivalent there.

Every returning path is an obstruction witness.

The intermediate states may reveal how the conflict arises.

A single bridge state may connect two otherwise separate local object systems.

Several independently valid correspondences may combine into a cycle.

A long chain may propagate identity through many contexts before returning to its starting domain.

The obstruction is not located at any one step. Every edge in the chain is locally licensed. The failure lies in their composition.

This distinguishes gluing obstruction from direct disagreement.

In direct disagreement, one observer identifies x and y while another observer sharing both states distinguishes them.

In a gluing obstruction, no single local identity declaration need conflict with another. The contradiction appears only after transitivity combines them.

The path structure can be represented by a graph G.

The vertices are states in X.

An edge connects two states whenever some Ei identifies them.

The connected components of G are the E*-classes.

For a fixed domain Ui, color each vertex by its Ei-class whenever it lies in Ui. Faithful descent fails exactly when one connected component contains two different colors belonging to Ui.

An obstruction witness is then a path inside that component joining vertices with different Ui-colors.

This graphical interpretation makes the obstruction concrete. It also shows that failure can spread.

Suppose one erroneous edge joins two large connected components. The resulting global class may contain many local objects from several domains. A single bridge can therefore create numerous defect pairs simultaneously.

The number of defect pairs may be large even when the number of causally decisive edges is small.

For diagnosis, one should therefore distinguish the visible extent of the defect from its minimal generating cause.

6.2 New Local Relations Generated Globally

The restriction E*|Ui is an equivalence relation on Ui. It is coarser than Ei because it contains Ei and may contain additional pairs.

The extra pairs are not added independently. They occur in blocks determined by mergers of Ei-classes.

Let

$Q_i = U_i/E_i$

be the local quotient.

The global relation E^* induces an equivalence relation $\approx_i$ on Q_i by

$A \approx_i B$

exactly when there exist $x \in A$ and $y \in B$ such that $x \, E^* \, y$.

Because every member of an E_i-class is already E^*-equivalent, the choice of representatives does not matter.

The relation $\approx_i$ records which local objects become fused under globalization.

Faithful descent on U_i occurs exactly when $\approx_i$ is equality on Q_i.

If A and B are distinct local objects but $A \approx_i B$, then every state in A becomes E^*-equivalent to every state in B. The global restriction merges the entire classes, not merely the original witness pair.

Thus the local defect is class-structured.

Suppose a global class C intersects U_i in the union

$C \cap U_i = A_1 \cup A_2 \cup \cdots \cup A_m$

of distinct E_i-classes.

Then globalization replaces those m local objects with one object inside the restricted global policy.

The defect relation contributed by C consists of all cross-class pairs

$(A_p \times A_q) \cup (A_q \times A_p)$

for $p \neq q$.

The total defect D_i is the union of these cross-class products over all global classes.

This representation separates two kinds of size.

The number of fused local classes measures ontological loss.

The number of defect pairs measures relational loss.

A merger of two very large classes may produce many defect pairs while reducing the local object count by only one.

Conversely, a global class containing many singleton local objects may create comparatively few pairs per merger but eliminate many distinct object identities.

The induced partition of Qi is therefore often more informative than Di alone.

Let

mi = |Qi|

when Ui is finite.

Let

m*i = |Ui/(E*|Ui)|.

Then

m*i ≤ mi.

The quantity

Li = mi − m*i

measures the number of local object distinctions lost under globalization.

Faithful descent on Ui is equivalent to

Li = 0.

The ratio

ρi = m*i / mi

may be called the local fidelity ratio.

It equals 1 under faithful descent and decreases as more local objects are fused.

These measures capture the consequence of obstruction but not its path complexity. Two networks can have the same local fidelity ratio while one fails through a single short bridge and the other through many long distributed chains.

A full obstruction analysis must therefore combine partition loss with witness structure.

6.3 The Gluing Obstruction Relation

For each local domain Ui, define the gluing obstruction relation

Oi = E*|Ui \ Ei.

The notation Oi and Di may be used interchangeably when emphasizing obstruction or defect, respectively.

The obstruction relation has several properties.

It is irreflexive because every diagonal pair belongs to Ei.

It is symmetric because both E*|Ui and Ei are symmetric. If (x, y) belongs to Oi, then (y, x) also belongs to Oi.

It is not generally transitive.

Suppose x and y belong to distinct Ei-classes fused globally, and y and z belong to another pair of distinct fused classes. It may happen that x and z are already Ei-equivalent, in which case (x, z) is not in Oi. More generally, set difference of equivalence relations need not be an equivalence relation.

The obstruction relation is therefore not itself a new identity policy. It is a diagnostic relation marking where the generated policy exceeds the local policy.

A domain is unobstructed when

Oi = ∅.

The full local system is unobstructed when

Oi = ∅

for all i.

This is equivalent to faithful descent.

The global obstruction set may be defined as the indexed family

O = {Oi}i∈I.

Indexing is necessary because the same pair of states may be considered under different local policies. A pair may be forbidden in one domain and absent from another domain entirely.

One may also define the set of obstructed domains:

Iobs = {i ∈ I : Oi ≠ ∅}.

This identifies where the generated global ontology fails to reproduce local ontology.

The obstruction can be expressed at the quotient level as

OQi = {(A, B) ∈ Qi × Qi : A ≠ B and ϕi(A) = ϕi(B)},

where

$\phi_i: Q_i \to X/E^*$

is the canonical local-to-global map.

Thus OQ_i is the kernel pair of ϕ_i with the diagonal removed.

The map ϕ_i is injective exactly when OQ_i is empty.

This quotient-level obstruction avoids counting every state pair and identifies the actual local objects that are fused.

For large or continuous state spaces, the quotient-level form may be mathematically preferable.

For finite computational systems, both forms are useful.

The state-level defect indicates affected records, detections, samples, or configurations.

The quotient-level defect indicates affected objects or categories.

The obstruction relation is canonical relative to the supplied local policies. It does not depend on which faithful globalization is attempted, because every extending globalization contains E^*.

If a pair belongs to O_i, it is identified by every global equivalence relation containing all local E_i.

The obstruction is therefore unavoidable unless the input system is altered.

This gives the obstruction relation a universal meaning:

O_i is the set of locally forbidden identities forced by every global acceptance of the local identity declarations.

6.4 Quantifying Descent Defects in Finite Systems

When X and the cover are finite, gluing defects can be measured exactly.

Several quantities capture different aspects of failure.

The first is the defect-pair count:

$P_i = |O_i|$.

Because O_i is symmetric, one may instead count unordered pairs:

$P_{ui} = P_i / 2$.

This measure records how many locally distinguished state pairs become globally identified.

The second is object-loss:

$L_i = |U_i/E_i| - |U_i/(E^*|U_i)|$.

This records how many local quotient objects disappear through fusion.

The third is the number of corrupted global traces.

Let C_i be the number of E^*-classes whose intersection with U_i contains more than one E_i-class.

Each such global class returns to U_i in a locally nonfaithful manner.

The fourth is maximum fusion multiplicity.

For every E^*-class C, let

$\mu_i(C)$

be the number of E_i-classes contained in $C \cap U_i$.

Define

$M_i = \max_C \mu_i(C)$.

Faithful descent gives $M_i = 1$ whenever U_i is nonempty.

A value $M_i = r$ means that at least one global object fuses r locally distinct objects in U_i.

The fifth is the mean fusion multiplicity over corrupted traces:

Averaging $\mu_i(C)$ over those C for which $\mu_i(C) > 1$ measures how broadly each defective global object collapses local ontology.

The sixth is obstruction depth.

For a defect pair $(x, y) \in O_i$, define

$\delta_i(x, y)$

as the length of the shortest local identity zig-zag connecting x and y.

Because x and y are not E_i-equivalent, $\delta_i(x, y)$ cannot be realized entirely as one E_i-step. It is at least two under the usual edge-count convention.

The minimum defect depth

$\delta_{min,i} = \min\{\delta_i(x, y) : (x, y) \in O_i\}$

identifies the shortest obstruction.

The maximum defect depth records the most distributed shortest witness among defect pairs.

A short depth suggests that the obstruction is locally concentrated.

A long depth suggests that the conflict emerges only through extended identity propagation.

The seventh measure is domain span.

For a witness path p, let σ(p) be the number of distinct observational domains used by its edges.

For a defect pair, minimize σ over all witness paths.

This gives the smallest number of observers required to generate that forbidden fusion.

The eighth measure is bridge sensitivity.

For each direct identity edge e, remove e and recompute the defect. The decrease in total defect indicates how much of the obstruction depends on that edge.

An edge whose removal eliminates many defect pairs is a high-impact bridge.

This measure is especially useful in entity matching and sensor association, where a single uncertain link may cause a large global cluster merger.

The ninth is domain responsibility, although the term must be interpreted cautiously.

For each domain Uj, remove all nonreflexive relations contributed uniquely by Ej and recompute the obstruction. The change estimates the structural contribution of that domain to global failure.

It does not establish that the observer is erroneous. The domain may contain valid information that merely conflicts with the intended global policy.

A normalized global defect score can be defined by combining local pair counts:

$$D_{pair} = \Sigma_i w_i \, |O_i| \, / \, |U_i \times U_i \setminus E_i|,$$

where w_i are selected weights and the denominator counts locally distinguished ordered pairs.

A quotient-based global score can be defined as

$$D_{obj} = \Sigma_i w_i \, L_i \, / \, |U_i/E_i|.$$

These summaries are useful for comparison but should not replace the full obstruction structure. Two systems with equal scores may differ greatly in which domains, objects, and paths are affected.

The most informative finite diagnosis should therefore report:

the obstructed domains,

the fused local classes,

minimal witness paths,

high-impact bridge edges,

and the resulting changes in local quotient cardinality.

6.5 Obstruction Witnesses and Minimal Zig-Zags

Every defect pair has at least one finite witness.

Let (x, y) ∈ Oi.

Then x E* y, so there exists a finite zig-zag

x = x0, x1, …, xn = y

with

xk−1 Eik xk

for every k.

Because x is not Ei-equivalent to y, the path cannot be reduced to an identity sequence contained entirely within one Ei-class.

A minimal witness is a witness with the smallest possible number of edges.

Minimal witnesses have useful structural properties.

A minimal witness contains no repeated state, except possibly when studying cycles explicitly. If a state repeats, the segment between its two appearances can be removed to produce a shorter path.

Thus a minimal witness is a simple path in the state identity graph.

A minimal witness cannot contain two consecutive steps that can be replaced by one direct local identity edge. If

xk−1 Ej xk

and

xk Ej xk+1

within the same local relation Ej, transitivity gives

xk−1 Ej xk+1,

allowing the intermediate state to be removed.

Therefore consecutive edges in a minimal witness generally represent meaningful changes of local class or observational domain.

At the quotient level, a minimal witness becomes a simple path through local object nodes. It begins at one Ei-class and ends at another distinct Ei-class belonging to the same domain.

The path may be viewed as an obstruction cycle after adding the fact that the two endpoint classes are locally separated. The network supports a route of positive identifications around the cycle, while the local policy imposes a negative separation across its endpoints.

Equivalence relations record positive sameness directly. Distinction is encoded by absence. For diagnostic purposes, it is useful to represent local distinction as an explicit negative edge.

Then an obstruction witness consists of:

a path of positive identity edges,

together with one negative local edge joining its endpoints.

This signed-cycle representation makes the inconsistency visible.

Every edge in the positive path demands fusion.

The negative edge forbids it.

A shortest signed cycle is a minimal obstruction certificate.

This structure resembles inconsistency detection in constraint systems. Positive constraints require elements to occupy one equivalence class. Negative constraints require selected pairs to occupy different classes.

The local observation system can therefore be reformulated as a constraint problem:

For every x Ei y, impose x = y at the object level.

For every x, y ∈ Ui with x not Ei-equivalent to y, impose x ≠ y within any faithful globalization.

Faithful descent exists exactly when the equality constraints do not imply the negation of any inequality constraint.

The generated relation E* computes the closure of the equality constraints.

The obstruction relation records violated inequality constraints.

This formulation is computationally valuable. A union-find structure can process positive equivalences. After all unions, each local distinction can be checked to determine whether its endpoints have entered the same component.

Explicitly enumerating every negative pair may be expensive. It is often enough to verify that no component contains more than one local equivalence class from the same domain.

Minimal witness recovery requires storing a spanning forest or provenance graph for the union operations. When a forbidden fusion is detected, the stored paths can reconstruct the chain responsible.

A strong computational implementation should not merely return “descent failed.” It should produce a certificate:

the affected domain Ui,

the distinct local classes A and B,

a representative $x \in A$,

a representative $y \in B$,

and a shortest or near-shortest zig-zag connecting them.

Such certificates convert abstract failure into interpretable evidence.

6.6 Local Consistency Without Global Fidelity

One of the most important consequences of the theory is that every local test can succeed while global descent fails.

Each Ei may be a valid equivalence relation.

Every pairwise overlap may satisfy exact agreement.

Every triple and higher overlap may inherit one common restricted relation.

Every local quotient may support stable transformations.

No observer need make an internally contradictory judgment.

Nevertheless, the combined system may generate an unfaithful global equivalence.

This phenomenon may be called globally induced inconsistency, although the local relations themselves remain logically consistent. The inconsistency lies between the total set of positive identity declarations and the total set of locally certified distinctions.

A simple pattern produces the failure.

Let A and B be distinct Ei-classes in one domain Ui.

Suppose A connects through a chain of overlap correspondences to local objects in other domains and eventually returns to B.

Each correspondence is locally valid.

The resulting connected component contains both A and B.

Global equivalence therefore merges them.

The network has formed a loop of identity around a locally maintained distinction.

This is not detected by checking only adjacent observers.

It requires path-sensitive analysis.

The phenomenon has analogues in many fields.

In distributed databases, every pairwise record match may exceed its confidence threshold, yet the transitive entity cluster may contain two records known locally to be different people.

In sensor networks, every handoff between adjacent sensors may be plausible, yet the completed track may merge two objects simultaneously visible to one sensor.

In biological taxonomy, every neighboring population comparison may support continuity, yet the resulting chain may join terminal populations treated as distinct under one regional criterion.

In semantic integration, every translation between neighboring vocabularies may be defensible, yet composing translations can identify concepts that one source language sharply distinguishes.

In scientific theory integration, every local correspondence may preserve a selected structure, yet the composite correspondence may erase a distinction necessary in the original theory.

These cases demonstrate that consistency is not purely local when identity is transitive.

The mathematical lesson is general:

Local agreement controls edges.

Global fidelity depends on paths.

This also explains why increasing the number of local correspondences can reduce fidelity. More overlap information may improve connectivity but also create additional routes by which identity returns destructively to a domain.

Distributed observation therefore faces a genuine tradeoff.

Too little gluing leaves one object fragmented across local systems.

Too much or insufficiently constrained gluing produces false fusion.

The correct objective is not maximal connectivity. It is maximal connectivity subject to every required local distinction.

6.7 Repair by Refinement, Separation, or Cover Modification

When faithful descent fails, the local system must be altered if an exact global ontology is still desired. Because every extending globalization contains E*, choosing a different global equivalence relation cannot eliminate the obstruction.

Repair must change the generating data or the fidelity requirements.

Several repair strategies are available.

The first is edge deletion.

Remove one or more local identity declarations so that the offending zig-zags no longer exist. In applications, this corresponds to rejecting uncertain record matches, sensor handoffs, class correspondences, or translation links.

The objective may be to remove a minimum number of edges or a minimum total confidence weight while eliminating all forbidden local fusions.

This becomes a constrained graph-cut or correlation-clustering problem.

The second is policy refinement.

Replace one or more Ei with finer equivalence relations that identify fewer states.

Refinement removes local positive identity edges. It can break global paths and restore descent.

However, refinement may fragment local objects that were previously useful. It increases resolution and reduces compression.

The third is policy coarsening.

Replace a local Ei with the coarser restriction E*|Ui.

This removes the obstruction by accepting the globally generated fusion.

Coarsening does not change E*. It changes the local policy so that the global identity is no longer considered forbidden.

This repair is appropriate when the global network has revealed a legitimate identity that the local observer could not detect.

It is inappropriate when the generated merger is erroneous or destructive.

The fourth is explicit separation constraints.

Retain local distinctions as active global constraints and reject any proposed identity edge whose closure would violate them.

Instead of accepting all local equivalences and testing afterward, the system incrementally admits identities only when they preserve current separations.

This produces an admissibility-controlled globalization rather than unconditional equivalence closure.

The fifth is cover enlargement.

Add a domain containing states whose relationship was previously mediated only through a long chain. The new observer can compare them directly and impose a policy that either validates or blocks the proposed identity.

Cover enlargement increases direct observability.

It can convert an underdetermined or hidden obstruction into an explicit local decision.

The sixth is cover refinement.

Split a problematic domain into smaller domains so that the two endpoints of a forbidden fusion no longer occur together.

This may make a globalization formally faithful, because the local distinction is no longer jointly stated in any one domain.

Such a repair should be treated cautiously. It can hide rather than resolve a conflict.

Cover refinement changes the standard of fidelity by removing comparisons.

The seventh is domain duplication or contextual separation.

A shared state may play different roles in different contexts. Treating it as one literal element of X may create an unwanted bridge. Replacing it with context-indexed copies can prevent automatic identity propagation.

The copies may then be related by explicit comparison morphisms rather than literal equality.

This strategy leads toward fibred categories, stacks, and contextual identity systems.

The eighth is policy stratification.

Instead of demanding one global equivalence relation, maintain several global policies indexed by task, scale, or observer class.

A pair may be equivalent for tracking but distinct for legal identity.

It may be equivalent thermodynamically but distinct microscopically.

It may be equivalent for diagnosis but distinct for treatment response.

Many apparent obstructions arise because heterogeneous purposes are forced into one undifferentiated identity policy.

The ninth is confidence-weighted or graded repair.

Replace exact equivalence edges with weighted associations and require a separate threshold or optimization procedure before global fusion.

This abandons the purely exact theory but may better represent noisy observation.

The tenth is local-policy revision after global evidence.

When a global path is judged scientifically valid, update Ei so that the formerly distinct endpoint classes become one. The revised local ontology incorporates the discovery made by distributed integration.

Repair is therefore not always deletion. It may be learning.

The correct repair depends on whether the obstruction reveals error, ambiguity, contextual plurality, or genuine emergence.

The formal theory identifies the available mathematical operations:

delete identities,

refine policies,

coarsen policies,

add separations,

modify domains,

duplicate contexts,

or replace one global ontology with a policy-indexed family.

6.8 An Obstruction Theory for Conceptual Integration

The preceding constructions support a general obstruction theory for distributed concept formation.

The input consists of local domains Ui and local equivalence relations Ei.

The generated global relation E* contains every identity implied by the local system.

The obstruction family

Oi = E*|Ui \ Ei

records every locally forbidden consequence.

A nonempty Oi certifies that no global equivalence relation can contain all local identities while preserving all local distinctions.

Each defect pair possesses a finite zig-zag witness.

Each witness can be reduced to a simple path.

At the quotient level, the obstruction appears as multiple local objects from one domain entering the same global component.

The theory therefore provides four equivalent levels of diagnosis:

relation-level excess,

quotient-level fusion,

graph-level component collision,

and path-level witness.

An obstruction analysis may proceed in the following order.

Construct or represent each local quotient Ui/Ei.

Create identity links induced by shared states or other licensed local correspondences.

Compute the connected global components.

For each domain, determine whether a component contains more than one local class.

Record every such collision.

Recover minimal paths connecting the collided classes.

Measure the size, depth, and bridge sensitivity of the resulting defects.

Classify candidate repairs according to their effect on local resolution and global connectivity.

This procedure turns conceptual integration into a certifiable mathematical operation.

A successful integration is accompanied by a proof of faithful descent.

A failed integration is accompanied by an obstruction certificate.

A repaired integration can be compared with the original system in terms of identities removed, distinctions added, objects split, or objects fused.

The obstruction theory also clarifies the relation between discovery and contradiction.

A generated identity that violates a local policy has two possible broad interpretations.

The global construction may be wrong relative to the intended ontology.

Or the local ontology may be incomplete relative to the global evidence.

Mathematically, both cases have the same initial signature:

$O_i \neq \emptyset$.

The difference emerges only after an admissibility judgment identifies which constraints are authoritative.

If local distinction is authoritative, the global path must be broken.

If the global identity is authoritative, the local policy must be coarsened.

If both are authoritative in different contexts, identity must become policy-indexed rather than globally absolute.

The obstruction is therefore not merely a defect to eliminate. It is a location where the system must decide what kind of objecthood it intends to preserve.

This gives gluing obstruction a role in discovery calculus.

A nonempty defect can reveal:

an erroneous correspondence,

a missing distinction,

an overly rigid local category,

an illegitimate transitive assumption,

a contextual dependence of identity,

or a new global object not visible locally.

The obstruction relation marks the exact boundary at which existing object systems cease to coexist under one global equivalence.

That boundary is mathematically informative.

It identifies where a conceptual framework must split, refine, fuse, or become indexed by context.

The resulting sequence is:

local identities generate a global closure,

global closure returns consequences to local domains,

local domains test those consequences against preserved distinctions,

and the resulting defects expose the points at which conceptual integration requires revision.

The next section develops the interpretive taxonomy of these outcomes. It distinguishes faithful gluing, direct gluing failure, false gluing, and emergent global identity, showing that the same formal closure mechanism can produce either destructive fusion or legitimate new objecthood.

7. False Gluing, Failed Gluing, and Emergent Identity

Equivalence closure is mathematically unambiguous but interpretively incomplete. Once local identity relations have been specified, the generated global equivalence E* is determined. The closure process does not distinguish whether a newly produced global identification is desirable, erroneous, destructive, or revelatory. It records only what follows when all local declarations of sameness are accepted and transitivity is enforced.

The same formal mechanism can therefore support radically different outcomes.

A chain of local identifications may correctly assemble partial observations of one persistent object.

A direct disagreement on an overlap may make exact integration impossible before globalization even begins.

A globally generated chain may merge objects that one local observer correctly distinguishes, producing false gluing.

Alternatively, that same type of chain may reveal that the local distinction was provisional and that a larger object exists only at the distributed level.

The mathematics of closure alone cannot decide among these interpretations because the distinction depends on what identity is intended to preserve. A relation appropriate for one purpose may be destructive for another. Two records may represent the same patient for epidemiological counting but require separation as different hospital encounters. Two physical states may represent one thermodynamic macrostate while remaining distinct microscopic states. Two organisms may belong to one evolutionary lineage while remaining distinct ecological types. Two visual detections may represent one tracked object while remaining separate images.

Objecthood is therefore indexed by an admissibility policy, transformation regime, explanatory task, or scale.

This section develops a taxonomy of gluing outcomes. It distinguishes faithful gluing, failed gluing through direct incompatibility, false gluing through excessive closure, and emergent identity generated by distributed evidence. It then examines destructive fusion, the difference between emergence and error, and the ontological status of objects that exist only through global integration.

The central conclusion is that a gluing obstruction is not yet a verdict. It is a demand for policy clarification.

7.1 Three Distinct Outcomes of Globalization

At the most basic level, a distributed observation system can produce three mathematically distinct outcomes.

The first is faithful gluing.

The local equivalence relations admit a global equivalence relation whose restriction to each domain is exactly the original local policy. Equivalently,

$E^* | U_i = E_i$

for every i.

No local object is fused with another object from the same domain. Every global equivalence class intersects each observational domain in at most one local equivalence class.

The second is failed gluing through incompatibility.

Two local policies disagree directly on an overlap. There exist i, j and states $x, y \in U_i \cap U_j$ such that

$x \; E_i \; y$

but

x is not Ej-equivalent to y.

No global equivalence relation can restrict exactly to both policies. The conflict is visible locally and requires no transitive closure to discover.

The third is unfaithful globalization through closure.

The local policies agree on all overlaps, but the generated global relation produces a returning zig-zag that merges locally distinct objects. There exists some i and x, y ∈ Ui such that

x E* y

but

x is not Ei-equivalent to y.

This third outcome contains two conceptually different cases.

The generated identity may be false relative to the intended object system. The integration has destroyed a distinction that should have been preserved.

Or the generated identity may be evidentially valid, revealing that the local policy should be revised. The integration has discovered a larger object not visible from the local domain alone.

The formal defect relation is the same in both cases.

Interpretation separates false gluing from emergent identity.

This produces a four-part conceptual taxonomy:

faithful gluing,

failed gluing,

false gluing,

emergent gluing.

The first three terms describe the relation between globalization and the original local policies.

The fourth describes a justified revision of those policies in response to the global construction.

The distinction between false gluing and emergent gluing cannot be made without an external criterion of admissibility. Such a criterion may involve prediction, causal structure, persistence, transformation invariance, intervention, explanatory compression, or task performance.

A global identification becomes legitimate only when the differences it erases remain irrelevant under the operations that matter.

Thus the taxonomy of gluing outcomes is inseparable from stable quotient theory. Descent determines whether local identities are preserved. Stability determines whether the resulting global identities support lawful transformation.

A global object should therefore satisfy two independent tests:

descent legitimacy,

dynamical legitimacy.

A quotient can be dynamically stable while unfaithful to a local ontology.

A quotient can descend faithfully while being unstable under the transformations required by the application.

Lawful distributed objecthood requires both.

7.2 Faithful Gluing

Faithful gluing occurs when a compatible family of local policies descends exactly to a global policy.

Let

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$.

The family glues faithfully when

$E^*|U_i = E_i$

for every i.

The corresponding global quotient

X/E^*

contains the finest global object system forced by the local identities.

Each local quotient map

$q_i: U_i \to U_i/E_i$

factors through the global quotient by an injective map

$\phi_i: U_i/E_i \to X/E^*$.

The map ϕi assigns each local object to the global object containing it. Injectivity guarantees that two distinct local objects from the same observer remain distinct globally.

Faithful gluing does not require local quotient spaces to remain separate. Objects belonging to different domains may be identified when overlap chains show that they are partial views of one global object.

Suppose camera A observes one segment of a trajectory, camera B observes the next segment, and camera C observes a third. The three local tracks may glue into one global track. Faithfulness requires only that no camera's two distinct local tracks be merged by the global chain.

Suppose an anatomical structure extends through many image slices. Each slice contains a local tissue region. Those regions may glue into one three-dimensional object. Faithfulness requires that the global reconstruction not merge two regions that one slice distinguishes as separate anatomical structures.

Suppose scientific databases contain partial records of one entity. Faithful integration may unify records from different databases while preserving every explicit within-database distinction.

The defining feature of faithful gluing is therefore not absence of fusion. It is controlled fusion.

Fusion across local contexts is permitted.

Fusion against local identity policy is not.

A faithfully glued global object may be distributed. No single domain need contain all its states. The object may exist only through a network of overlapping local classes.

Let C be an E*-class. For every i,

C ∩ Ui

is either empty or exactly one Ei-class.

The collection of these nonempty traces forms a compatible family of local objects. The global class is their union together with the identity paths connecting them.

Faithful gluing therefore turns local objects into sections or traces of larger distributed objects.

The process preserves local ontology while extending its range.

It is useful to distinguish strong and weak forms of faithful gluing.

Strong faithful gluing occurs when the global equivalence relation is uniquely determined by the local data. The minimal globalization E* is then the only faithful global relation.

Weak faithful gluing occurs when E* is faithful but additional faithful globalizations exist. The local data determine a finest global ontology but leave some cross-domain identities unresolved.

In the weak case, states belonging to disconnected or insufficiently compared regions may be merged or separated without changing any local restriction. The observation network contains underdetermined identity.

Faithful gluing therefore does not always imply complete global knowledge.

It guarantees consistency, not uniqueness.

A distributed system may be globally coherent while still leaving open whether certain remote objects are identical.

This distinction is important in science. Absence of local contradiction does not amount to proof of global uniqueness. Additional observation may be required to resolve underdetermined object correspondences.

Faithful gluing supplies a coherent baseline from which such further inquiry can proceed.

7.3 Failed Gluing Through Incompatibility

Failed gluing occurs when local policies disagree directly on a common domain.

Let

$U_{ij} = U_i \cap U_j$.

If

$E_i|U_{ij} \neq E_j|U_{ij}$,

then there exist shared states x and y about which the two observers make different identity judgments.

One observer may identify them.

The other may distinguish them.

No global equivalence relation can recover both policies exactly.

This failure is immediate. It does not arise from long-range transitivity, global closure, or returning paths. It is already present in the local data.

A direct incompatibility can take several forms.

The observers may use different measurement resolutions.

One observer may classify two states as the same because it records only a coarse feature. Another may distinguish them using additional information.

The observers may serve different tasks.

A medical billing system may treat two records as separate encounters, while a clinical identity system treats them as one patient.

The observers may use different temporal assumptions.

One database may regard two names as belonging to one continuously existing entity, while another treats a legal or organizational change as the creation of a new entity.

The observers may use different causal criteria.

Two phenomena may be grouped by appearance in one discipline and separated by mechanism in another.

The observers may simply contain error.

A sensor may misidentify an object, a database may contain a duplicate, or a classifier may produce a false match.

Mathematically, all of these appear as disagreement between restricted equivalence relations.

There are several possible responses.

One policy may be declared authoritative.

A common refinement may preserve every distinction made by either observer:

$Ei|Uij \cap Ej|Uij$.

This resolves conflict by treating a pair as identical only when both observers agree.

A common coarsening may accept every identity declared by either observer:

$EqCl(Ei|Uij \cup Ej|Uij)$.

This resolves conflict by treating a pair as identical whenever either observer does, together with all transitive consequences.

A new negotiated policy may be introduced.

The contexts may remain separate under policy-indexed identity.

Or the overlap states may be duplicated into context-specific copies, preventing literal conflict while preserving comparison data.

Each resolution changes the original descent problem.

Exact faithful gluing of the unrevised policies remains impossible.

This is why failed gluing should not be confused with false gluing.

Failed gluing concerns incompatible inputs.

False gluing concerns compatible inputs whose global closure produces an unacceptable result.

In failed gluing, the contradiction is visible at an interface.

In false gluing, every interface may appear consistent.

The distinction matters computationally.

Direct incompatibility can be detected by checking overlap restrictions.

False gluing requires global connected-component or path analysis.

The distinction also matters institutionally.

A direct disagreement can be brought to the observers for explicit negotiation.

A false gluing may remain hidden until a large integrated system produces an anomalous merger.

The latter is often more dangerous because no individual local decision appears incorrect.

7.4 False Gluing Through Excessive Closure

False gluing occurs when equivalence closure merges states or local objects that should remain distinct under the intended global policy.

Formally, let E^* be the minimal globalization. A local defect occurs when

$E_i \subset E^*|U_i$.

The mathematical system has accepted every local equivalence declaration and enforced transitivity. The result is logically correct relative to those positive declarations. It is false only relative to an additional requirement that certain distinctions remain globally authoritative.

False gluing is therefore not a failure of equivalence closure.

It is a failure of the input identity policy or of the assumption that all local equivalences may be composed without qualification.

A common cause is an erroneous bridge.

Suppose two large identity components are correctly formed. A single mistaken association connects one state in the first component to one state in the second. Transitive closure merges the entire components.

The global damage can be much larger than the local error.

In entity resolution, one false match may merge the records of two people, companies, publications, or devices.

In object tracking, one incorrect handoff may combine two trajectories.

In knowledge graphs, one mistaken equivalence assertion may collapse distinct concepts.

In biological classification, one overly permissive correspondence may merge lineages that require separation.

Another cause is context loss.

An identity relation may be valid only under one task or scale but may be exported as if it were unconditional.

For example, two materials may be equivalent for thermal modeling but not for structural engineering. If their local thermodynamic equivalence is used inside a global engineering ontology, transitive closure may erase mechanically essential differences.

A third cause is noncongruence of local policies.

Two local observers may each use valid equivalence relations, but the substitutions licensed by one context may not compose with those licensed by another.

Equivalence within each domain does not guarantee admissibility of cross-domain composition.

A fourth cause is shared-state ambiguity.

An overlap state may represent different roles, times, modes, or contextual occurrences. Treating the shared token as one literal state creates a bridge that should instead have been represented by a relation between context-indexed copies.

A fifth cause is semantic drift.

A chain of locally plausible correspondences can gradually change meaning. Each adjacent translation is acceptable, yet the endpoints differ substantially. Transitive identity then confuses relatedness with sameness.

False gluing exposes a limitation of exact equivalence as a universal integration mechanism.

Equivalence is transitive without attenuation. Once x is equivalent to y and y to z, the system must identify x and z completely. There is no memory of path length, confidence, contextual change, or semantic drift.

In noisy or heterogeneous systems, this can be too strong.

The appropriate repair may require:

typed equivalence,

graded similarity,

bounded path composition,

context-indexed identity,

provenance-sensitive closure,

or admissibility constraints on which relations may compose.

False gluing can be diagnosed through obstruction witnesses. A returning zig-zag connects two local classes that should remain separate.

The witness should then be examined for its weakest or least admissible step.

The problem may not be any individual edge considered locally. It may be the transition between relation types.

For example, morphological equivalence may compose with geographical equivalence and then with functional equivalence, but the composite may not preserve the property required by the original domain.

The future algebra of distributed concept formation must therefore classify not only states and equivalence relations but also types of identity evidence and admissible compositions among them.

False gluing is the phenomenon that makes this enrichment necessary.

7.5 Emergent Global Identity

Not every globally generated identity absent from a local policy is erroneous.

A distributed system may reveal a unity that no local observer could establish alone.

This is emergent global identity.

Suppose x and y lie in one domain Ui and are locally distinguished. A chain through other domains connects them. The resulting relation

$x \; E^* \; y$

creates a descent defect relative to Ei.

If the global evidence is accepted as authoritative, the correct response is not to break the chain. It is to revise Ei so that x and y become members of one broader local class.

From the original local perspective, descent failed.

From the revised global perspective, discovery occurred.

Emergent identity therefore appears mathematically as a justified coarsening of a local policy induced by distributed information.

Several examples illustrate the phenomenon.

A local astronomical observation may record two appearances as distinct sources. Measurements from other times or wavelengths may reveal that they are images of one gravitationally lensed object.

A laboratory may classify two samples as distinct strains. Genomic and epidemiological data from other institutions may show that they belong to one transmission lineage.

A local database may treat two author names as different people. Cross-database affiliations, coauthors, identifiers, and publication histories may demonstrate that they are one person.

A biological observer may distinguish life stages as different forms. Distributed developmental evidence may reveal them as phases of one organism.

A set of local coordinate descriptions may appear unrelated until transition data show that they describe one global geometric object.

A theory may treat several phenomena as separate until an invariant or symmetry reveals that they are manifestations of one underlying structure.

In these cases, global closure does not merely combine known objects. It changes what counts as an object.

Emergent identity is not arbitrary fusion. It must be certified.

Several criteria may support certification.

The global class may exhibit persistence under admissible transformations.

It may improve prediction.

It may unify causal history.

It may support interventions that act coherently across its local traces.

It may reduce explanatory complexity without destroying required distinctions.

It may possess invariants not visible in any one domain.

It may survive independent reconstruction from different subsets of the observer network.

These criteria move beyond descent into stable objecthood.

A global identity is emergent when the generated equivalence produces a quotient object that is more explanatory, predictive, or transformation-stable than the separated local objects, while the erased local distinction is shown to be nonessential for the intended purpose.

Emergence therefore requires a change of policy.

Let Ei be the original local relation and let

E'i = E*|Ui.

Replacing Ei with E'i accepts every global identification that returns to Ui.

The revised family is automatically faithful relative to E*, because the local policies are now defined as restrictions of the global relation.

However, indiscriminate replacement of all Ei by E*|Ui merely forces consistency. It does not prove scientific validity.

The revision becomes meaningful only when the global quotient is independently certified.

Emergent identity is thus a two-stage process:

global closure proposes a new object,

admissibility and stability tests certify it.

The closure generates a hypothesis of objecthood.

The transformation regime decides whether that hypothesis should be retained.

7.6 Destructive Fusion of Locally Distinct Objects

Destructive fusion is the loss of a distinction required for correct reasoning, prediction, action, or explanation.

It is the operational consequence of false gluing.

Suppose local objects A and B in Ui are merged into one global class C. If some admissible transformation, measurement, intervention, or judgment treats A and B differently, then the global quotient cannot support that operation representative-independently.

There exist $x \in A$ and $y \in B$ such that $x \; E^* \; y$, but for some relevant transformation g,

$g(x)$ is not E^*-equivalent to $g(y)$,

or some relevant observable h satisfies

$h(x) \neq h(y)$.

The quotient then erases information required by the task.

Destructive fusion can therefore be detected as instability.

This connects gluing defects with the destructive abstraction theorem from stable quotient theory.

A global identity relation may be coherent as an equivalence relation while failing to be invariant under the transformations that matter.

The resulting global objects exist set-theoretically but not operationally.

Destructive fusion may take several forms.

Predictive destruction occurs when fused states have different future outcomes.

Causal destruction occurs when fused states respond differently to intervention.

Semantic destruction occurs when fused concepts support different inferential roles.

Legal destruction occurs when merged records refer to entities whose obligations or rights must remain separate.

Biological destruction occurs when merged classes differ in reproduction, pathology, development, or ecological function.

Computational destruction occurs when abstracted program states permit different transitions or outputs.

Physical destruction occurs when a coarse identification erases variables required for the dynamics under study.

The seriousness of a fusion depends on the intended application.

A distinction may be safely erased for one operation but destructive for another.

Two patient records may be fused for identifying a person but must remain distinct for tracking separate clinical events.

Two molecular configurations may be fused for equilibrium thermodynamics but separated for reaction kinetics.

Two words may be fused as synonyms in one context but distinguished in legal or technical use.

The appropriate response is often not to reject global identity absolutely but to index it by policy.

Let π denote a task or admissibility regime.

Then identity should be written conceptually as

$x \equiv_\pi y$.

The same states may satisfy one policy-relative equivalence and fail another.

Destructive fusion occurs when an equivalence certified under π is used under a different policy σ without verifying that the required distinctions remain irrelevant.

This suggests that distributed gluing should eventually be performed in a category of policy-indexed equivalences rather than in one undifferentiated lattice Eq(X).

Each local identity edge should carry information about what substitutions it licenses.

Composition should be allowed only when those licenses remain compatible.

Until that enriched theory is developed, transformation testing supplies the main defense against destructive fusion.

A global object must not merely be connected.

It must behave as one object under the operations for which it is introduced.

7.7 Distinguishing Emergence from Observational Error

Emergence and error can produce the same initial mathematical pattern.

A local policy distinguishes x and y.

Distributed closure identifies them.

The defect relation records the pair.

Nothing in the relation E* alone determines whether the local distinction should be preserved or abandoned.

A principled distinction requires additional evidence.

Several tests can be applied.

The first is independent path support.

If the global identity depends on one fragile edge, error is more plausible. If several structurally independent chains connect the same local objects, the identity receives stronger support.

Independence must be defined relative to the application. Paths sharing the same underlying data source or algorithm may not provide independent evidence.

The second is perturbation stability.

Remove uncertain edges, vary thresholds, change the cover, or perturb measurements. A legitimate emergent object should persist across reasonable changes. A false fusion may disappear when one weak link is removed.

The third is transformation invariance.

Apply the admissible dynamics. If different representatives of the proposed global object continue to behave coherently, the identity is supported. If they diverge in task-relevant ways, the fusion is destructive.

The fourth is predictive improvement.

Does the global object enable more accurate prediction than the separated local objects? Improvement should be evaluated without concealing failures in important subgroups or contexts.

The fifth is causal unity.

Do the local traces share one causal history, mechanism, or intervention response? Mere association may support connectivity without supporting identity.

The sixth is explanatory compression.

Does the global class reduce description length or theoretical complexity while preserving the phenomena that require explanation? Compression that discards anomalies rather than explaining them is not sufficient.

The seventh is recurrence across representation.

Can the same global object be reconstructed from different measurements, feature sets, languages, or observer subsets? Representation-independent recovery supports emergence.

The eighth is boundary coherence.

Does the proposed object have a stable boundary against neighboring global classes? A component created through diffuse weak associations may lack a meaningful identity even if internally connected.

The ninth is reversibility of refinement.

If the global class is later refined, do the local objects reappear as coherent subobjects? Legitimate emergence often preserves lower-level structure as organized internal differentiation rather than obliterating it.

The tenth is intervention success.

Can one act on the proposed global object as one unit and obtain coherent effects? Operational unity is strong evidence of objecthood.

No single test is universally decisive.

The appropriate standard depends on the domain.

The tests together establish an important methodological principle:

Connectivity proposes identity; invariance and intervention certify identity.

Observation networks are good at generating candidate objects because closure discovers connected structures. They are not sufficient by themselves to establish that every connected component is one lawful object.

The distinction between emergence and error should therefore be represented as a model-selection problem over identity policies.

Let E* be the generated proposal.

Let T be the family of admissible transformations.

Let L measure predictive or explanatory loss.

Let C measure complexity or compression.

A candidate global policy may be evaluated by whether it:

contains sufficiently supported local identities,

respects authoritative distinctions,

is invariant under T,

minimizes L,

and achieves useful compression C.

The exact optimization depends on the application, but the structural requirement is general.

Object formation is a constrained selection among possible quotient policies, not merely an automatic consequence of connectedness.

The defect relation identifies where the selection decision must be made.

7.8 The Ontology of Globally Generated Objects

A globally generated object is an equivalence class of E*.

Its unity may not be directly observed from any single domain. It exists as the closure of a distributed system of local identifications.

What kind of object is this?

It is not necessarily a fundamental object independent of all observation.

It is not merely a subjective label chosen without constraint.

It is an object relative to an organized network of identity policies.

Its members are connected by licensed local substitutions.

Its global unity is the transitive consequence of those substitutions.

Its legitimacy depends on whether the resulting class remains stable under the transformations, predictions, explanations, and interventions that matter.

This position avoids two extremes.

The first extreme treats only directly and globally observable entities as real. Under that view, distributed objects are regarded as constructions lacking genuine ontology.

The second extreme treats every connected closure as automatically real. Under that view, any network of local identifications creates an object regardless of stability or purpose.

The present theory supports neither.

A distributed object can be real at an effective level even when no observer sees it whole. Hurricanes, ecosystems, economies, languages, species, organizations, trajectories, and scientific theories are often known through partial observations. Their objecthood depends on coherent persistence and lawful behavior, not on simultaneous total visibility.

At the same time, distributed construction is fallible. Identity networks can contain erroneous edges, inappropriate compositions, contextual shifts, and destructive abstractions.

The ontology of a generated object is therefore certified rather than presumed.

A useful hierarchy can be stated.

A candidate object is a connected component of local identity declarations.

A coherent object is a candidate whose globalization satisfies faithful descent or whose local policy revisions are explicitly accepted.

A stable object is a coherent object whose equivalence class is preserved by admissible transformations.

An explanatory object is a stable object that supports useful prediction, compression, intervention, or unification.

An autonomous effective object is an explanatory object whose behavior can be described at the quotient level without continuous reference to all underlying states.

This hierarchy separates mere closure from mature objecthood.

The theory also permits genuinely relational objects.

A global object need not possess one privileged center, representative, or intrinsic description. It may be constituted by the pattern of equivalences among local presentations.

Its identity lies in the coherence of its traces.

This is closely related to the categorical principle that an object may be characterized by its relations and transformations rather than by an inaccessible internal essence.

Distributed objecthood extends that principle.

The object is represented by a compatible family of local manifestations and the descent structure joining them.

When strict descent fails but higher coherence remains meaningful, the correct ontology may no longer be one equivalence class in a set. It may require a groupoid, stack, homotopy type, or other higher object that retains multiple local presentations and the transformations among them.

False gluing can therefore signal that the chosen ontological level is too coarse.

Two local objects may not be literally identical, yet they may be equivalent through a nontrivial path that should not be collapsed.

Replacing all equivalence paths by equality destroys information about how the identifications were established.

A higher theory would retain:

the local objects,

the equivalences between them,

the alternative paths of equivalence,

and the relations among those paths.

The set-theoretic quotient forgets this structure.

For many applications, that forgetting is appropriate.

For others, it creates destructive fusion.

The distinction points toward the next levels of the research program.

Exact set-based descent supplies the foundational case.

Policy-indexed identity supplies contextual control.

Enriched morphisms record types and costs of identification.

Groupoids retain multiple presentations.

Stacks organize local objects that glue only up to equivalence.

Higher categories retain transformations among gluing procedures.

The present section establishes why those extensions are needed.

Faithful gluing shows that local object systems can sometimes form one coherent global ontology without modification.

Failed gluing shows that direct disagreement prevents exact integration.

False gluing shows that compatibility alone does not protect against destructive transitive closure.

Emergent identity shows that a descent defect may also be the mathematical signature of discovery.

The final distinction depends on what the global object preserves.

A generated identity is false when it destroys structure required by the intended theory.

It is emergent when it reveals a larger stable structure for which the former local distinction is no longer essential.

The local-to-global calculus can therefore be summarized as follows:

agreement licenses comparison,

closure generates candidate unity,

restriction reveals conflict,

stability tests objecthood,

and admissibility decides whether the conflict should be repaired by separation or by conceptual revision.

The next section adds dynamics to this structure. It asks when locally admissible transformations assemble into global transformations and when those transformations preserve the generated equivalence, allowing globally glued objects to persist and evolve lawfully.

8. Stable Dynamics on Glued Quotient Spaces

A global object system is not mathematically complete merely because its local identities can be glued. Objecthood also requires persistence under transformation. The resulting quotient classes must support lawful dynamics that do not depend on the choice of representative.

This requirement was central to the single-domain stable-quotient theory. If E is an equivalence relation on a state space X and g: X → X is an admissible transformation, then g descends to the quotient X/E only when

x E y implies g(x) E g(y).

The condition ensures that equivalent representatives remain equivalent after transformation. The induced quotient map

$\tilde{g}([x]E) = [g(x)]E$

is then well defined.

Distributed observation introduces additional layers of coherence.

A local observer may possess a transformation defined only on its own domain. Different observers may describe different parts of the same global evolution. Their local transformations must agree where their domains overlap, or else they cannot assemble into one global transformation. Even if they agree as functions, they must preserve the local identity policies. Finally, their assembled global action must preserve the equivalence relation generated by all local identities.

The dynamics of distributed object formation therefore involve three distinct questions.

Can the local transformations be assembled into a global transformation?

Do the local transformations preserve the local equivalence relations?

Does the assembled transformation preserve the generated global equivalence?

Under strict assumptions, the first two conditions imply the third. If locally equivalent states are sent to locally equivalent states and local maps assemble into one global function, every global zig-zag is transformed into another global zig-zag. The generated equivalence is therefore invariant.

This gives a stable globalization theorem.

The theorem shows that local dynamical coherence can descend to the global quotient just as local identity can descend to a global object system. But the two forms of descent are not identical.

Identity descent asks whether global equivalence preserves all original local distinctions.

Dynamical descent asks whether transformations preserve the global equivalence classes that have been formed.

A system can satisfy one condition and fail the other.

The generated quotient may be dynamically stable while unfaithful to one or more local policies.

A faithfully descended quotient may fail to support a proposed global transformation.

A complete distributed object system requires both faithful identity formation and stable transformation.

This section develops local transformation systems, overlap compatibility, global assembly, invariance of generated equivalence, induced quotient dynamics, and the distinction between local and global stability. It concludes by examining ways in which distributed abstractions fail dynamically even when their static gluing appears coherent.

8.1 Local Admissible Transformation Monoids

Let X be a state space covered by observational domains $\{U_i\}_{i \in I}$.

Each U_i carries a local observation policy E_i.

A local transformation on U_i is a function

$g: U_i \to U_i$

or, more generally, a function

$g: U_i \to U_j$

between observational domains.

The first case represents an internal transformation of one local context.

The second represents transport of states between contexts.

The initial theory will emphasize endomorphisms

$g: U_i \to U_i$

because they act repeatedly within one domain and form the local analogue of a dynamical system.

A collection G_i of local transformations on U_i is an admissible transformation monoid when it contains the identity map and is closed under composition.

Thus

$\mathrm{id}_{U_i} \in G_i$,

and if $g, h \in G_i$, then

$g \circ h \in G_i$.

The monoid structure represents repeatable local action.

Transformations may describe time evolution, computational updates, symmetry operations, coordinate changes, perturbations, interventions, or other processes regarded as lawful within the local theory.

The observation policy E_i is stable under G_i when every $g \in G_i$ preserves E_i:

$x \, E_i \, y$ implies $g(x) \, E_i \, g(y)$.

The pair

(U_i, E_i, G_i)

is then a local stable quotient system.

Every $g \in G_i$ induces a unique quotient transformation

$\bar{g}: U_i/E_i \to U_i/E_i$

defined by

$\bar{g}([x]Ei) = [g(x)]Ei$.

The induced transformations form a monoid $\bar{G}i$ acting on the local quotient.

The quotient map

$qi: Ui \rightarrow Ui/Ei$

is equivariant:

$qi \circ g = \bar{g} \circ qi$.

This means that transforming a state and then passing to its object class gives the same result as first passing to the class and then applying the induced object-level transformation.

The local quotient therefore supports autonomous dynamics.

The observer can reason in terms of local objects without selecting or tracking particular representatives.

Local admissibility is policy-relative.

A transformation may preserve one equivalence relation while failing to preserve another.

Suppose Ei groups images by object identity while ignoring illumination. A brightness transformation may preserve Ei. A transformation that replaces one depicted object with another does not.

Suppose Ei groups physical microstates by total energy. Hamiltonian evolution may preserve the equivalence if energy is conserved, while external forcing may not.

Suppose Ei groups program states by final output. An execution step may preserve the eventual-output equivalence if termination behavior is unchanged, but a memory corruption may not.

The selection of Gi therefore specifies what the objects are expected to survive.

A quotient class is not stable absolutely. It is stable under a designated family of transformations.

In a distributed system, each domain may carry a different monoid Gi. These monoids may represent different local descriptions of one global process, or they may represent distinct transformation regimes.

For the local transformations to participate in a common global dynamics, additional compatibility conditions are required.

8.2 Compatibility of Transformations on Overlaps

Suppose Ui and Uj overlap.

Let gi: Ui → Ui and gj: Uj → Uj be local transformations.

The most direct compatibility condition is equality on the overlap:

gi(x) = gj(x)

for every x ∈ Ui ∩ Uj.

However, this equation is meaningful only when both images lie in a common codomain and the transformations preserve the overlap or are regarded as restrictions of one global map.

A strict overlap-preserving condition is

gi(Ui ∩ Uj) ⊆ Ui ∩ Uj

and

gj(Ui ∩ Uj) ⊆ Ui ∩ Uj,

together with

gi|Ui∩Uj = gj|Ui∩Uj.

Under this condition, the two observers assign the same transformed state to every state they both observe.

This is functional compatibility.

A weaker condition may be sufficient when exact state agreement is unnecessary. The observers may produce states that differ but remain observationally equivalent.

One may then require

gi(x) Eij gj(x)

for every x ∈ Ui ∩ Uj,

where Eij is the common restricted equivalence relation on the overlap.

This is quotient-level compatibility.

The strict set-theoretic assembly theorem uses functional compatibility. Quotient-level compatibility belongs to an enriched theory in which local maps assemble only up to observational equivalence.

The present section begins with strict compatibility.

Let {gi}i∈I be a family of local transformations.

The family is overlap-compatible when

gi(x) = gj(x)

for every x ∈ Ui ∩ Uj.

This condition guarantees that the rule

g(x) = gi(x) whenever x ∈ Ui

does not depend on the selected domain.

Because the domains cover X, the rule defines a function on all of X.

The resulting function is the assembled global transformation.

Overlap compatibility is necessary for exact assembly.

If x lies in Ui ∩ Uj and gi(x) ≠ gj(x), there is no single-valued global function whose restrictions are both gi and gj.

The conflict resembles failed gluing of identity policies. It is a direct disagreement on shared data.

Transformation compatibility must also account for domains after the transformation.

If each gi maps Ui into itself, the assembled g maps every x ∈ Ui into Ui. A state belonging to several domains may remain in each only if the local maps agree and each preserves the relevant domain.

More general dynamics may move states between domains. A sensor may hand an object to a neighboring sensor. A time evolution may carry a state out of one coordinate patch. A computational process may migrate data between machines.

Such cases require local maps

gij: Uij-source → Uj

or a category of domains and transitions rather than independent endomorphism monoids.

The strict endomorphism framework is therefore foundational rather than exhaustive.

It captures distributed transformations that can be represented as restrictions of one global endomorphism.

Within this framework, overlap compatibility has a second level.

The induced local quotient transformations should agree on shared object traces.

Suppose the local policies agree on $U_i \cap U_j$. If g_i and g_j agree pointwise on the overlap and preserve E_i and E_j, then their induced quotient actions agree on the overlap quotient.

Thus strict state-level compatibility implies quotient-level compatibility.

The converse need not hold. Two transformations may send a shared state to different representatives of the same quotient object.

A future weak-descent theory should permit such transformations to glue up to equivalence rather than equality.

That extension is one of the routes from sets toward groupoids and higher categorical observation.

8.3 Global Transformations Assembled from Local Actions

Global Assembly Proposition.

Let $\{U_i\}_{i \in I}$ be a cover of X. For every i, let

$g_i: U_i \to X$

be a function.

Assume that for all i and j,

$g_i|_{U_i \cap U_j} = g_j|_{U_i \cap U_j}$.

Then there exists a unique function

$g: X \to X$

such that

$g|_{U_i} = g_i$

for every i.

Proof.

For $x \in X$, choose an index i with $x \in U_i$ and define

$g(x) = g_i(x)$.

If x also belongs to U_j, overlap compatibility gives

$g_i(x) = g_j(x)$.

Thus g is well defined.

It restricts to each g_i by construction.

If $h: X \to X$ has the same restrictions, then for every x choose U_i containing x. One has

$h(x) = g_i(x) = g(x)$.

Therefore $h = g$.

This proves existence and uniqueness.

The proposition is the ordinary gluing law for functions. Functions are determined pointwise, so agreement on overlaps is sufficient.

This contrasts sharply with equivalence relations. Compatible local equivalence relations may fail faithful descent because transitivity generates nonlocal consequences. Compatible local functions do not have this problem. Their global assembly requires no additional closure.

This difference is structurally important.

Transformation gluing is pointwise.

Identity gluing is closure-sensitive.

When each g_i maps U_i into U_i, the assembled global map g preserves every domain:

$g(U_i) \subseteq U_i$.

If the family of local transformations is indexed by an element a of a common monoid G, write

$g_{ai}: U_i \to U_i$

for the local action of a on U_i.

Assume

$g_{a\,i}|_{U_i \cap U_j} = g_{a\,j}|_{U_i \cap U_j}$

for every $a \in G$ and every i, j.

Then each a assembles into a global transformation

$g_a: X \to X$.

If the local actions satisfy

$g_{ei} = id_{U_i}$

and

gabi = gai ∘ gbi

for all a, b ∈ G, their assembled maps satisfy the same laws globally:

ge = idX,

gab = ga ∘ gb.

Thus G acts on X.

The proof follows by restriction. Both sides agree on every Ui and hence agree globally.

A family of compatible local monoid actions therefore assembles into a global monoid action.

One may also begin with different local monoids Gi. To construct one global action, the local monoids must be coordinated by a common indexing monoid, by compatible homomorphisms, or by a category of local transformations.

Without such coordination, there is no canonical meaning to applying “the same” transformation across all domains.

For example, each sensor may possess its own clock and update operation. A global time evolution exists only after their updates are synchronized or related through explicit transition data.

Each scientific model may possess its own admissible operations. A global transformation theory requires a rule for translating operations between models.

The algebra of distributed dynamics therefore contains two gluing problems:

gluing transformations as functions,

and gluing the operation systems that index them.

The present theorem assumes the latter problem has been solved by a common transformation monoid or by a specified compatible family.

8.4 Invariance of Generated Equivalence

Suppose the local maps assemble into a global transformation g: X → X.

Assume each local policy Ei is preserved by the corresponding restriction gi:

x Ei y implies gi(x) Ei gi(y).

Because g|Ui = gi, this can be written as

x Ei y implies g(x) Ei g(y).

The question is whether g preserves the generated global equivalence

E* = EqCl(⋃i Ei).

The answer is yes.

Generated Invariance Proposition.

If g: X → X preserves every local equivalence relation Ei, then g preserves E*.

That is,

x E* y implies g(x) E* g(y).

Proof.

Suppose x E* y.

By the zig-zag characterization, there exists a finite sequence

x = x0, x1, …, xn = y

such that for each k,

xk−1 Eik xk

for some local relation Eik.

Because g preserves every local relation,

g(xk−1) Eik g(xk).

Therefore

g(x0), g(x1), …, g(xn)

is another local identity zig-zag.

Its endpoints are g(x) and g(y).

Hence

g(x) E* g(y).

This proves invariance.

The proposition shows that equivalence closure is functorial with respect to maps preserving the generating relations.

A transformation need not preserve every local domain individually for the path argument to work. It is enough that every generating relation is mapped into the generated global relation:

x Ei y implies g(x) E* g(y).

This weaker condition permits a local equivalence pair to move into another domain or into a relation supported only globally.

Thus two levels of invariance can be distinguished.

Strong local invariance requires

x Ei y implies g(x) Ei g(y).

Weak generated invariance requires

x Ei y implies g(x) E* g(y).

Strong local invariance immediately gives weak generated invariance.

Weak generated invariance is sufficient for descent to X/E*.

A transformation can therefore be lawful globally even if it does not preserve each local quotient internally.

It may transport one local object into a distributed global object whose traces cross several domains.

This distinction becomes important in systems where objects move between observational regions.

The minimal condition for quotient dynamics is simply

x E* y implies g(x) E* g(y).

Because E* is generated by local edges, it is sufficient to verify the condition on those generators:

For every i,

x Ei y implies g(x) E* g(y).

This generator test can be more efficient than checking all pairs in E*.

In categorical language, g is a morphism from the policy-indexed space (X, E*) to itself.

In computational language, g maps every connected component of the local identity graph into one connected component.

It may send different global objects to the same object, but it may not split one global object into several images depending on representative.

The image of an E*-class must lie entirely within one E*-class.

This is the operational meaning of generated invariance.

8.5 The Stable Globalization Theorem

The preceding results combine into the principal dynamical theorem.

Stable Globalization Theorem.

Let X be a set covered by domains {Ui}i∈I. Let Ei be an equivalence relation on Ui, and define

E* = EqCl(⋃i Ei).

Let G be a monoid. Suppose that for every a ∈ G and every i, there is a local map

gai: Ui → Ui

satisfying the following conditions.

First, the local maps define an action on each domain:

gei = idUi,

gabi = gai ∘ gbi.

Second, the local maps agree on overlaps:

gai|Ui∩Uj = ga j|Ui∩Uj.

Third, every local action preserves the local policy:

x Ei y implies gai(x) Ei gai(y).

Then:

For every a ∈ G, the family {gai} assembles uniquely into a global map

ga: X → X.

The maps {ga} define an action of G on X.

Every ga preserves E*.

Therefore every ga induces a unique map

ḡa: X/E* → X/E*

given by

$\bar{g}a([x]E^*) = [ga(x)]E^*$.

The induced maps define an action of G on X/E^*.

The global quotient map

$q^*: X \to X/E^*$

is equivariant:

$q^* \circ ga = \bar{g}a \circ q^*$.

Proof.

Overlap agreement gives a unique global map ga for every a.

The monoid action laws hold because they hold on every Ui. For $x \in X$, choose Ui containing x. Then

$ge(x) = gei(x) = x$,

and

$gab(x) = gabi(x) = gai(gbi(x)) = ga(gb(x))$.

Thus G acts on X.

Each ga preserves every Ei by assumption. The generated invariance proposition therefore implies that ga preserves E^*.

The induced map $\bar{g}a$ is well defined.

The identity and composition laws descend:

$\bar{g}e([x]) = [ge(x)] = [x]$,

and

$\bar{g}ab([x]) = [gab(x)] = [ga(gb(x))] = \bar{g}a(\bar{g}b([x]))$.

Thus G acts on X/E^*.

Equivariance follows directly:

$(q^* \circ ga)(x) = [ga(x)]E^* = \bar{g}a([x]E^*) = (\bar{g}a \circ q^*)(x)$.

This proves the theorem.

The stable globalization theorem establishes representative-independent dynamics on the minimally generated global objects.

It does not require faithful identity descent.

Even if E* fuses local objects, the global quotient can still support a lawful action.

This is why stability and fidelity must be evaluated separately.

The theorem also has a weaker version.

Suppose the local transformations assemble into global maps ga, but local invariance holds only into E*:

x Ei y implies ga(x) E* ga(y).

Then the same zig-zag proof shows that ga preserves E*. The quotient action still exists.

This generalized form is more appropriate when local objects move between domains.

The strict theorem is cleaner and supplies the foundational case.

8.6 Induced Dynamics on the Global Quotient

Once E* is invariant under the global transformation monoid G, the quotient

X/E*

becomes a dynamical system.

Each global object [x]E* has an image

$\bar{g}([x]E^*) = [g(x)]E^*$.

The image does not depend on which representative x is chosen.

This allows the dynamics to be studied entirely at the object level.

Several structural phenomena become available.

A global object C may be fixed:

$\bar{g}(C) = C$.

This means that g maps every state in C back into C.

The internal states may change while the object identity persists.

A collection of global objects may form an orbit under G.

A transformation may permute global objects.

Several global objects may converge to one object under a noninvertible action.

One global object cannot split into different quotient images under a deterministic induced map. If different representatives would lead to different objects, the quotient action would not be well defined.

Thus representative-independent dynamics prohibit object-level branching caused solely by hidden choice of representative.

Branching may still occur under nondeterministic transformations, relations, or stochastic kernels. Those require a broader framework.

The quotient action can reveal effective dynamics not obvious at the state level.

A large number of microscopic transitions may induce a simple transition among a few quotient classes.

A distributed trajectory reconstructed from many sensor states may become one path through global object space.

A network of local classifications may induce transitions among globally integrated concepts.

A scientific theory may track changes in effective objects rather than all underlying descriptions.

The degree of compression is determined by E*, while the validity of the compressed dynamics is determined by invariance.

If X is finite, the quotient dynamics can be represented by a directed graph on E*-classes.

For each class C, draw an arrow to $\bar{g}(C)$.

For a monoid with several generators, label arrows by the corresponding transformations.

This global quotient graph may be much smaller than the original state-transition graph.

The reduction is exact with respect to the selected equivalence.

No information relevant to the induced quotient dynamics is lost, because all representatives of one class have quotient-equivalent images.

The quotient may also possess invariants inherited from the global action.

A function

$F: X/E^* \to Y$

defines an E*-invariant state-level function

$F \circ q^*: X \to Y$.

Conversely, any function $h: X \to Y$ constant on E^*-classes factors uniquely through X/E^*.

If h is also invariant under G, the induced quotient function is invariant under the quotient action.

Thus distributed object formation and dynamical invariance interact through factorization.

The global quotient is the natural domain for every observable that respects all locally generated identities.

The quotient dynamics describe how those observables can evolve without recovering discarded local detail.

8.7 Local Stability Versus Global Stability

Local stability and global stability are related but not identical.

Local stability means that each transformation preserves each local equivalence relation.

Global stability means that the assembled transformation preserves E^*.

Strong local stability implies global stability.

The converse does not generally hold.

A global transformation may preserve E^* while failing to preserve a particular E_i.

Suppose x and y are E_i-equivalent. Their images may cease to be E_i-equivalent but remain E^*-equivalent through other domains. The global object persists even though its local presentation changes.

This can occur when an object moves across observational regions.

Two states initially represented as one local object may evolve into states whose unity is visible only through a different observer or through a distributed path.

Global identity remains coherent, but local identity is not invariant.

This suggests two types of object persistence.

Local persistence means that an object remains identifiable within the same local quotient.

Distributed persistence means that its global identity survives even when its local traces change, disappear, or move between domains.

A moving physical object illustrates the distinction. It may leave one camera's field of view and enter another. The first local track ends, but the global trajectory continues.

A biological lineage may disappear from one region while persisting elsewhere.

A semantic concept may change vocabulary while retaining a global inferential role.

A mathematical object may move between coordinate charts while remaining one point or geometric structure globally.

Local instability need not imply global extinction.

Conversely, every local quotient may be dynamically stable while the global assembly fails.

Several causes are possible.

The local transformations may disagree on overlaps, so no global function exists.

They may assemble into a global function that maps locally equivalent pairs consistently but fails to preserve cross-domain correspondences used in E*.

The local monoids may represent incompatible notions of time or action.

A local object correspondence may be valid statically but not respected dynamically.

For example, two local tracks may be matched at one time, but their subsequent evolutions may diverge in ways incompatible with one global trajectory.

Thus static overlap identity does not automatically imply dynamical compatibility.

A complete local-to-global theory should track both local equivalence classes and the transformations connecting their temporal traces.

One may define four regimes.

Local and global stability both hold.

The local objects persist, the global object system persists, and quotient dynamics are coherent at every level.

Local stability holds but global stability fails.

Each observer has a valid internal dynamics, but the local systems do not assemble into one lawful global evolution.

Global stability holds but some local stability fails.

The global objects persist, but their local presentations change or migrate across domains.

Neither local nor global stability holds.

The selected policies fail to support the proposed transformation regime at either level.

These regimes help separate modeling errors from genuine movement between contexts.

The strongest form of distributed objecthood occurs when faithful descent and stable globalization both hold.

Then:

local identities glue without distortion,

local transformations assemble,

the global equivalence is invariant,

and the quotient supports representative-independent dynamics.

This may be called a faithfully stable distributed quotient.

8.8 Dynamical Failure of Distributed Abstractions

A distributed abstraction fails dynamically when its global object classes do not support the transformations required by the theory.

Let E^* be the generated global relation.

Suppose there exists a transformation $g: X \to X$ and states x, y such that

$x \; E^* \; y$

but

$g(x)$ is not E^*-equivalent to $g(y)$.

Then g does not descend to X/E^*.

The class $[x]E^*$ has no unique image. Choosing x as representative produces one quotient object, while choosing y produces another.

The global object is dynamically ambiguous.

Several forms of failure can produce this condition.

The first is hidden local instability.

A local policy identifies states that evolve into locally or globally distinct outcomes. The defect already exists in one domain and propagates to the global quotient.

The second is unstable bridge identity.

Two local objects are joined through an overlap correspondence that is valid at one moment but not preserved by evolution.

The bridge creates one global class, but the later states separate.

The third is asynchronous transformation.

Local observers apply updates at different times or under different clocks. Their maps agree statically but fail to represent one global step.

The fourth is incompatible intervention.

A transformation applied in one domain has no corresponding action in another, or it changes the meaning of shared states.

The fifth is contextual drift.

The equivalence relations themselves change over time. States that counted as the same under Ei at time t may be distinguished under a revised policy at time t + 1.

A fixed quotient E* cannot represent this policy evolution without extension.

The sixth is object bifurcation.

One effective object develops into several distinguishable descendants. Deterministic quotient dynamics can represent one class mapping to one later class only if the later identity policy still groups those descendants. To represent genuine fission, the framework must permit time-indexed equivalence relations or relational dynamics between quotient spaces.

The seventh is object fusion.

Several classes may map into one class without violating deterministic quotient dynamics. Noninvertible quotient maps naturally represent fusion. But if the transformation is expected to be reversible, such fusion reveals loss of information incompatible with the intended dynamics.

The eighth is local-to-global path destruction.

A global object may be held together by a chain of local equivalences. A transformation can remove an intermediate bridge, causing the class to fragment under the newly generated relation.

If the policy E* is kept fixed, the old object persists formally despite loss of observational support.

If E* is recomputed from transformed local policies, the object may split.

This reveals a distinction between fixed-policy dynamics and endogenous policy dynamics.

Fixed-policy dynamics hold the equivalence relation constant and ask whether states transform compatibly with it.

Endogenous policy dynamics allow the identity policy itself to evolve as observations, domains, or admissibility criteria change.

The present theorem treats fixed policies.

A full discovery calculus must model trajectories

$E_t \rightarrow E_{t+1}$

in the space of global identity policies.

Object birth, fusion, fission, maintenance, and extinction then become changes in the quotient system itself.

Dynamical failure should therefore be diagnosed at two levels.

At the state level, ask whether g preserves E*.

At the policy level, ask whether the transformed observation network regenerates the same equivalence relation.

Let the local policies after transformation be E'i. Their globalization is

$E'^* = \mathrm{EqCl}(\bigcup_i E'_i)$.

If

$E'^* = E^*$,

the object ontology is dynamically stationary.

If E'* refines E*, global objects split or new distinctions arise.

If E'* coarsens E*, global objects fuse.

If neither relation refines the other, the ontology reorganizes in a mixed way.

This policy-dynamic perspective connects stable globalization with theory dynamics and concept formation.

A distributed object is maintained when its global class remains stable both under state transformation and under regeneration of the observational identity system.

A class may persist as a set-theoretic quotient while losing observational support.

It may remain observationally connected while failing transformation invariance.

Robust objecthood requires both.

The practical diagnosis of a distributed abstraction should therefore include the following tests.

Check whether local transformations preserve local equivalences.

Check whether local transformations agree on overlaps.

Check whether they assemble into a global action.

Check whether every generating identity edge maps into E*.

Check whether every global quotient class has a unique quotient image.

Check whether recomputed policies after transformation preserve the same object structure.

Check whether any global class relies on fragile bridge identities destroyed by small perturbations.

Check whether local-to-global object correspondences remain consistent over time.

When these conditions hold, the global quotient supports stable dynamics.

When they fail, the obstruction should be reported with the same precision used for static descent defects.

A dynamical obstruction certificate consists of:

two E*-equivalent states x and y,

an admissible transformation g,

and evidence that g(x) and g(y) lie in different E*-classes.

At the local-object level, it consists of one global object whose representatives acquire incompatible object-level images.

At the path level, it may identify the local equivalence edge or overlap correspondence that ceases to be preserved.

Such certificates expose exactly why a proposed global object cannot function as one object under the selected dynamics.

The combined static and dynamical theory can now be summarized.

Local policies generate local objects.

Compatible local policies generate candidate global objects.

Faithful descent determines whether local distinctions survive globalization.

Compatible local transformations assemble into global transformations.

Generated invariance determines whether the global objects survive those transformations.

The quotient supports lawful dynamics only when equivalent representatives remain equivalent after every admissible action.

The full sequence is therefore:

local identity,

overlap compatibility,

global closure,

faithful descent,

local transformation,

functional assembly,

global invariance,

quotient dynamics.

This sequence establishes the minimal mathematical conditions for a distributed observation network to produce not merely a static classification, but a persistent world of effective objects.

The next section reorganizes the local policy assignments categorically. It studies the contravariant assignment $U \mapsto Eq(U)$, identifies it as a presheaf of conceptual resolutions, and explains precisely why ordinary sheaf gluing fails for equivalence relations.

9. The Presheaf of Conceptual Resolutions

The local-to-global theory of object formation has so far been developed in set-theoretic terms. A state space X is covered by domains U. Each domain carries a lattice $Eq(U)$ of observation policies. Inclusion of a smaller domain into a larger one induces restriction of policies. Compatible families of local relations may or may not arise as exact restrictions of one global relation.

These constructions have a natural categorical organization.

The assignment

$U \mapsto \mathrm{Eq}(U)$

is contravariant. A larger observational domain carries policies that can be restricted to smaller domains. If V is contained in U, then every equivalence relation on U determines an equivalence relation on V. The direction of the policy map is therefore opposite to the direction of domain inclusion.

This is the defining pattern of a presheaf.

A presheaf assigns mathematical data to each domain and restriction maps to each inclusion, subject to identity and composition laws. The data may be functions, groups, rings, vector spaces, categories, or other structures. In the present case, the data are complete lattices of equivalence relations.

The resulting object may be called the presheaf of conceptual resolutions.

It records every possible identity policy available on every observational domain and the way those policies appear when attention is restricted.

The presheaf perspective clarifies several features of the theory.

First, conceptual resolution is localizable. One may ask what distinctions remain visible in any subdomain.

Second, local policies can be compared through common restrictions.

Third, a global policy determines a compatible family of local policies automatically.

Fourth, the converse is not guaranteed. A compatible family may fail to glue faithfully because transitive closure creates additional local identifications.

This fourth point separates the presheaf from a sheaf.

A sheaf is a presheaf in which compatible local data can be glued uniquely into global data. Equivalence relations fail this property under unrestricted covers. Exact overlap agreement does not ensure that a compatible local family is the restriction of one global equivalence relation.

The failure is not caused by poor notation or an incomplete choice of restriction maps. It is generated by the algebra of equivalence itself. Transitivity allows identities to travel outside a local domain and return with new consequences.

The presheaf therefore contains a precise obstruction to ordinary sheaf behavior.

This obstruction is mathematically useful. It identifies the difference between local consistency and global objecthood. It also shows that sheafification, if applied without care, may change the local policies rather than merely complete them. Enforcing global gluing can create new identities, merge local objects, and alter the intended ontology.

This section defines the presheaf formally, proves its elementary properties, describes the failure of the sheaf condition, examines separatedness and uniqueness, and interprets sheafification as a form of enforced conceptual closure.

9.1 Observation Policies Assigned to Domains

Let X be a fixed state space.

Consider a collection of observational domains contained in X. The collection may be the full power set of X, the open sets of a topology, the measurable subsets of a measurable space, the subgraphs of a graph, or another category of admissible domains.

For the elementary construction, let Dom(X) denote the category whose objects are subsets $U \subseteq X$ and whose morphisms are inclusions.

There is at most one morphism

$V \to U$

exactly when $V \subseteq U$.

Thus Dom(X) is a thin category determined by the inclusion order.

To each domain U assign the set

Eq(U)

of all equivalence relations on U.

This set is not merely a collection. It is a complete lattice under refinement.

For $E, F \in Eq(U)$, define

$E \preceq F$

when

$E \subseteq F$.

The relation E is then finer than F.

The meet of a family $\{E_\alpha\}$ is

$\bigwedge_\alpha E_\alpha = \bigcap_\alpha E_\alpha$.

The join is

$\bigvee_\alpha E_\alpha = \mathrm{EqCl}(\bigcup_\alpha E_\alpha)$.

The bottom element is equality:

$\Delta U = \{(x, x) : x \in U\}$.

The top element is universal equivalence:

$\nabla U = U \times U$.

Thus every domain is assigned a complete lattice of possible object-forming policies.

This assignment has a direct interpretation.

Eq(U) is the space of all possible conceptual resolutions on U.

Each point of Eq(U) specifies what counts as one object within U.

Moving downward in refinement preserves more distinctions.

Moving upward fuses more states into common objects.

The lattice therefore organizes all possible local ontologies available on the same underlying domain.

One may regard Eq(U) as a category as well as a lattice. There is a unique arrow

$E \to F$

when $E \preceq F$.

The arrow represents lawful coarsening or forgetting of distinctions.

Thus the value assigned to U can be treated as a thin category Res(U) of conceptual resolutions.

The presheaf may therefore be regarded either as lattice-valued or category-valued.

The lattice-valued view emphasizes meets, joins, and order.

The category-valued view emphasizes restriction functors and later generalization toward stacks.

For the present section, both interpretations will be used.

9.2 Contravariant Restriction Along Inclusions

Let $V \subseteq U$.

Every policy $E \in Eq(U)$ restricts to an equivalence relation on V:

$res_{U,V}(E) = E|V = E \cap (V \times V)$.

This defines a map

$res_{U,V}: Eq(U) \to Eq(V)$.

The map is monotone.

If

$E \subseteq F$,

then

$E|V \subseteq F|V$.

Therefore refinement is preserved under restriction.

The direction is contravariant.

The inclusion is

$V \hookrightarrow U$,

while the restriction map is

$Eq(U) \to Eq(V)$.

A policy on the larger domain yields a policy on the smaller domain.

A policy on V does not generally determine a policy on U. Extension requires choices and may be impossible under additional constraints.

This reversal of direction is natural for observation.

To restrict the field of view is easy. One forgets states outside V and retains all identity judgments among states still visible.

To enlarge the field of view is not automatic. New states appear, and their identities relative to existing states must be decided.

Restriction therefore behaves like information loss.

Extension behaves like object formation.

The maps $res_{U,V}$ preserve arbitrary meets:

$res_{U,V}(\bigcap_\alpha E_\alpha) = \bigcap_\alpha res_{U,V}(E_\alpha)$.

This follows because restriction is intersection with $V \times V$.

They preserve the bottom and top relations:

$\mathrm{res}_{U,V}(\Delta_U) = \Delta_V$,

$\mathrm{res}_{U,V}(\nabla_U) = \nabla_V$.

They do not generally preserve joins.

For relations E and F on U, one may have

$\mathrm{res}_{U,V}(E \vee F) \neq \mathrm{res}_{U,V}(E) \vee \mathrm{res}_{U,V}(F)$.

The left side first generates global equivalence paths in U and then restricts their consequences to V.

The right side first deletes states outside V and then generates equivalence using only paths remaining inside V.

A path may leave V and return. Its endpoint relation appears on the left but not on the right.

One always has

$\mathrm{res}_{U,V}(E) \vee \mathrm{res}_{U,V}(F) \preceq \mathrm{res}_{U,V}(E \vee F)$.

The locally generated join is at least as fine as the restriction of the globally generated join.

Global closure can create additional local identifications.

This failure of join preservation is the lattice-theoretic signature of nonlocal identity propagation.

The restriction maps are therefore complete meet-homomorphisms but not generally complete lattice homomorphisms.

This asymmetry becomes central in the failure of sheaf gluing.

9.3 The Resolution Presheaf

Define a contravariant functor

$\mathrm{Eq}_X: \mathrm{Dom}(X)^{op} \to \mathrm{CLat}_{\wedge}$,

where $\mathrm{CLat}_{\wedge}$ denotes a category of complete lattices and meet-preserving monotone maps.

The functor is defined on objects by

$\mathrm{Eq}_X(U) = \mathrm{Eq}(U)$.

For an inclusion

$V \subseteq U$,

it is defined on morphisms by

$\mathrm{EqX}(V \hookrightarrow U) = \mathrm{res}_{U,V}$.

To verify that this assignment is a functor, two laws must be checked.

The first is identity.

For every U,

$\mathrm{res}_{U,U}(E) = E$

for all $E \in \mathrm{Eq}(U)$.

Therefore

$\mathrm{EqX}(\mathrm{id}_U) = \mathrm{id}_{\mathrm{Eq}(U)}$.

The second is composition.

If $W \subseteq V \subseteq U$, then

$\mathrm{res}_{V,W}(\mathrm{res}_{U,V}(E)) = (E|V)|W = E|W = \mathrm{res}_{U,W}(E)$.

Thus

$\mathrm{EqX}(W \hookrightarrow V) \circ \mathrm{EqX}(V \hookrightarrow U) = \mathrm{EqX}(W \hookrightarrow U)$.

The assignment is therefore a presheaf.

It may be called the resolution presheaf, the observation-policy presheaf, or the presheaf of conceptual resolutions.

A section of this presheaf over U is simply an equivalence relation on U.

A local section over U_i is a local observation policy E_i.

A family $\{E_i\}$ on a cover $\{U_i\}$ is compatible when the sections agree under restriction to every overlap:

$E_i|U_i \cap U_j = E_j|U_i \cap U_j$.

A global section $E \in \mathrm{Eq}(X)$ restricts to a compatible family automatically.

Thus there is a canonical map

$\mathrm{Eq}(X) \to \mathrm{Match}(\{U_i\})$,

where Match({Ui}) denotes the set of compatible families on the cover.

This map sends

$E \mapsto \{E|Ui\}$.

The sheaf question is whether this map is bijective for every cover.

Injectivity would mean that a global policy is uniquely determined by all its local restrictions.

Surjectivity would mean that every compatible family of local policies comes from a global policy.

For equivalence relations, both questions require care.

Surjectivity fails in general because of gluing obstruction.

Injectivity can also fail when the cover does not compare enough pairs to determine cross-domain identity.

Thus the resolution presheaf may fail both existence and uniqueness parts of the ordinary sheaf condition.

The two failures have different meanings.

Failure of surjectivity means compatible local ontologies cannot coexist exactly inside one global ontology.

Failure of injectivity means the local ontologies do not determine all global identities.

The first is obstruction.

The second is underdetermination.

The presheaf framework separates them cleanly.

9.4 Identity and Composition of Restriction

The restriction laws are elementary, but their conceptual interpretation is important.

The identity law states that observing a domain without changing the domain leaves the identity policy unchanged.

No new distinctions are created and none are forgotten.

The composition law states that successive restriction is path-independent.

If $W \subseteq V \subseteq U$, restricting a policy from U to V and then from V to W gives the same result as restricting directly from U to W.

Thus local identity does not depend on the sequence through which observational attention narrows.

This law supports coherent comparison across nested domains.

Suppose an instrument first selects a region V and then zooms to W. The identity relation visible in W is the same as if W had been selected directly.

Suppose a scientific database is filtered by country and then by institution. The induced identity policy on the final set is the restriction of the original relation, independent of the intermediate filter.

Suppose a mathematical chart is restricted to a smaller neighborhood. The resulting equality or equivalence judgments do not depend on the chain of subdomains.

These laws make restriction reliable.

They also reveal that the difficulty in object formation does not lie in moving from global to local.

The presheaf structure is exact in that direction.

The difficulty lies in reversing the process.

Restriction forgets information monotonically.

Gluing must reconstruct information while respecting transitivity.

The loss incurred by restriction can be analyzed.

A global policy E on U may restrict to the same local policy on V as another global policy F even when $E \neq F$.

This occurs when their differences involve pairs not jointly visible in V.

Thus restriction is not generally injective.

A family of restrictions over a cover may collectively recover E if the cover is sufficiently informative.

If every pair $x, y \in U$ occurs together in some cover element, then the family of local restrictions determines E uniquely. For any pair, one reads its identity status from a domain containing both.

If some pairs never occur together, their global relation may remain underdetermined.

The identity and composition laws guarantee coherence of restriction, but not completeness of local recovery.

That is the role of separatedness.

9.5 Failure of the Ordinary Sheaf Condition

A presheaf is a sheaf when every compatible family of local sections has a unique global section whose restrictions reproduce the family.

For the resolution presheaf, this would require:

Given a cover {Ui} of U and compatible equivalence relations Ei on each Ui, there exists a unique equivalence relation E on U such that

E|Ui = Ei

for every i.

This condition fails in general.

Existence can fail.

Consider the domains

U1 = {a, b, d},

U2 = {b, c},

U3 = {c, d}.

Let E1 identify a with b but distinguish d.

Let E2 identify b with c.

Let E3 identify c with d.

The pairwise overlaps are singletons, so the local relations agree on every overlap.

The family is compatible.

Any global equivalence containing the local identities must identify

a, b, c, d

through transitivity.

Its restriction to U1 therefore identifies a with d, contradicting E1.

No global section exists.

Thus the presheaf fails the existence part of the sheaf axiom.

Uniqueness can also fail.

Let X be the disjoint union of U1 and U2.

Choose any equivalence relation E1 on U1 and E2 on U2.

The local family is compatible because the overlap is empty.

One global relation may keep every E1-class distinct from every E2-class.

Another may identify selected classes across the two domains.

Both restrict to the same local policies.

Thus several global sections exist.

The presheaf fails uniqueness.

The ordinary sheaf condition therefore fails in two independent ways.

Compatible local policies may be overconstrained globally because transitive closure creates forbidden identities.

Or they may be underconstrained globally because the cover does not determine cross-domain identities.

The first failure may be called obstructed gluing.

The second may be called nonunique gluing.

The map

Eq(U) → Match({Ui})

can fail to be surjective or injective.

The faithful descent theorem characterizes surjectivity for one compatible family. The generated relation E* provides the least possible global extension. The family is in the image exactly when E* restricts faithfully.

Injectivity depends on whether distinct global relations can have identical local traces.

A cover that contains every pair of states in at least one domain guarantees injectivity.

This property may be called pair-detecting coverage.

Ordinary topological covers rarely satisfy pair detection. Distant points need not lie in one open set of a selected cover. Consequently, global equivalence between distant local objects can remain invisible locally.

The resolution presheaf is therefore not generally a sheaf on the ordinary site of subsets with arbitrary covers.

This conclusion is not a defect of the theory.

It is one of its principal results.

Equivalence relations are intrinsically global structures because transitivity can connect arbitrarily distant states.

They cannot always be reconstructed from pairwise overlap agreement alone.

9.6 Separatedness and Uniqueness of Globalization

A presheaf is separated when compatible local data have at most one global gluing.

For the resolution presheaf, separatedness over a cover means:

If E and F are equivalence relations on U such that

$E|U_i = F|U_i$

for every i,

then

$E = F$.

This property depends on the cover.

A sufficient condition is pair detection.

Definition.

A cover $\{U_i\}$ of U is pair-detecting when for every $x, y \in U$, there exists an index i such that

$x, y \in U_i$.

Proposition.

If a cover is pair-detecting, then the resolution presheaf is separated over that cover.

Proof.

Let E and F have identical restrictions to every U_i.

Take any $x, y \in U$.

Choose U_i containing both states.

Then

x E y

if and only if

x E|Ui y.

Because the restrictions agree,

x E|Ui y

if and only if

x F|Ui y.

Hence

x E y

if and only if

x F y.

Therefore E = F.

Pair detection is sufficient but stronger than necessary.

Uniqueness may still hold when not every state pair lies directly in one domain, provided transitivity and the local policies determine the missing relations indirectly.

For a fixed compatible family {Ei}, uniqueness can be expressed at the level of minimal global classes.

Let E* be the minimal faithful globalization.

Another faithful relation F may merge distinct E*-classes only if doing so does not force two local classes from the same domain to merge.

Therefore E* is the unique faithful globalization exactly when every pair of distinct E*-classes is forbidden from merging by at least one local domain or by the transitive consequences of such constraints.

A simple sufficient condition is:

For every pair of distinct E*-classes C and D, there exists Ui meeting both C and D in locally distinct Ei-classes.

Then no faithful globalization can merge C and D.

The local system fully separates the minimal global objects.

When this condition fails, identity between some global components is underdetermined.

The set of faithful globalizations may then be studied as a partially ordered family above E*.

Its members represent alternative resolutions of identities not decided by local observation.

This family is closed under intersections but not generally under joins.

Thus even the ambiguity space of global ontology has nontrivial order structure.

Separatedness has an epistemic interpretation.

A separated observation network determines at most one global identity policy.

A nonseparated network permits multiple global ontologies that look identical from every local viewpoint.

No local observer can distinguish among them.

Additional domains, bridge measurements, or external principles are required.

This is a form of observational underdetermination.

Two global worlds may have the same local object structure but differ in which distant objects are identical.

The presheaf formalism makes that underdetermination exact.

9.7 Sheafification as Enforced Conceptual Closure

Every presheaf can often be associated with a sheaf through a process called sheafification. Informally, sheafification modifies the original presheaf so that compatible local data glue uniquely.

For conceptual resolutions, such a process must be interpreted cautiously.

A naive attempt to sheafify by taking compatible local equivalence relations and assigning their generated global closure

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$

does not preserve the original local sections when faithful descent fails.

The restriction $E^*|U_i$ may be coarser than E_i.

Thus the process does not merely glue the local policies. It revises them.

The revision identifies locally distinct objects whenever global transitivity requires it.

This may be called enforced conceptual closure.

Enforced closure takes a distributed family of identity declarations and completes it into a globally coherent equivalence relation, even when doing so changes the local ontology.

The resulting family

$\{E^*|U_i\}$

is compatible and globally realizable by construction.

It is the smallest globally induced coarsening of the original family.

For every i,

$E_i \subseteq E^*|U_i$.

Thus enforced closure preserves all positive local identity declarations while discarding some local distinctions.

This construction resembles a closure operator on families of local policies.

Let C denote the operation

$C(\{E_i\}) = \{E^*|U_i\}$.

Then C is extensive:

$E_i \subseteq C(\{E_i\})_i$.

It is monotone:

coarsening the input policies cannot produce a finer closure.

It is idempotent:

applying the operation to an already globally induced family changes nothing.

The fixed points are exactly the families arising as restrictions of their generated global equivalence.

These are effective descent families.

Thus enforced closure projects arbitrary local identity systems onto globally closed ones.

The operation has a clear mathematical role.

It computes the nearest family above the original policies, in the refinement order, that is induced by one global equivalence relation.

However, “nearest” here means least coarsening, not minimal semantic distortion under every possible measure.

A single required fusion can have extensive practical consequences.

The operation is therefore structurally canonical but not automatically normatively acceptable.

In some applications, enforced closure is desirable.

Entity records known to belong to one person should be merged consistently across every local database.

Coordinate presentations of one point should be recognized as one global point.

Distributed observations of one moving object should form one trajectory.

In other applications, enforced closure is destructive.

A semantic chain should not automatically collapse all related concepts into one meaning.

A chain of biological similarity should not automatically produce taxonomic identity.

A sequence of uncertain record matches should not automatically merge all records.

Sheafification in this setting is therefore not neutral completion.

It is an ontological policy.

It chooses global coherence over preservation of every local distinction.

The choice may be justified, but it must be explicit.

9.8 What Sheafification Preserves and Destroys

Enforced conceptual closure preserves several structures.

It preserves every locally declared equivalence.

If x Ei y, then x E* y, so the identity remains valid after closure.

It preserves global transitivity.

All chains of local identity are completed into coherent global classes.

It preserves minimality among globally coherent coarsenings.

No finer global equivalence relation can contain every local identity.

It preserves compatibility after revision.

The restricted family {E*|Ui} agrees on overlaps because every member is derived from the same global relation.

It also preserves any transformation that respects the generating local identities. Such transformations preserve E* and descend to the global quotient.

At the same time, enforced closure can destroy important structures.

It can destroy local distinctions.

Several Ei-classes may become one class under E*|Ui.

It can destroy injectivity of local-to-global object maps.

Distinct local objects may acquire the same global image.

It can destroy explanatory specificity.

A coarse global class may combine states requiring different local explanations.

It can destroy causal discrimination.

States with different intervention responses may be merged.

It can destroy provenance.

Once equivalence classes are collapsed, the particular path by which identity was established may be forgotten.

It can destroy contextual information.

An equivalence valid in one domain may propagate into domains where its justification does not apply.

It can destroy higher structure.

Multiple distinct equivalence paths become indistinguishable after set-theoretic quotienting.

The quotient records only that two states are equivalent, not how, through which observers, or with what evidence.

These losses suggest that ordinary sheafification may be too coarse for a mature theory of distributed observation.

Several alternatives are possible.

One may retain the full identity graph rather than only its connected components.

One may attach provenance to each equivalence path.

One may distinguish types of local equivalence and restrict which types may compose.

One may use groupoids so that equivalences remain morphisms rather than collapsing immediately into equality classes.

One may use stacks so that local objects glue up to specified equivalence rather than strict identity.

One may use higher categories to retain transformations among alternative gluing paths.

One may use graded or probabilistic structures when identity evidence has varying strength.

The need for these extensions is already visible in the presheaf failure.

The resolution presheaf shows that exact identity policies localize cleanly but do not always globalize faithfully.

The failure has three possible responses.

Reject some local identities so that faithful descent becomes possible.

Accept the generated coarsening and revise local ontologies.

Retain a richer object in which local distinctions and global equivalences coexist without being collapsed into one set-theoretic quotient.

The third response is the categorical direction developed in the next section.

The main conclusions of the presheaf analysis can now be summarized.

Each domain U carries a complete lattice Eq(U) of conceptual resolutions.

Every inclusion $V \subseteq U$ induces a contravariant restriction map.

These assignments form a presheaf of complete meet-lattices.

Global policies restrict to compatible local families.

Compatible local families need not possess global realizations.

Global realizations, when they exist, need not be unique.

Existence fails through transitive gluing obstruction.

Uniqueness fails through observational underdetermination.

Pair-detecting covers guarantee uniqueness but not existence.

Minimal globalization provides a canonical closure, but it may coarsen local policies.

Thus the presheaf is not an ordinary sheaf under unrestricted covers.

Its failure is itself the mathematics of distributed object formation.

A sheaf would say that compatible local objects determine one global object automatically.

The present theory says something more discriminating:

local compatibility is only evidence,

transitive closure generates a candidate global ontology,

faithful descent tests whether that ontology preserves local distinction,

and higher structures may be required when strict gluing is either impossible or destructive.

The next section expresses these ideas through categorical diagrams, descent objects, limits, and observer networks. It also prepares the transition from strict presheaves of equivalence relations toward stacks and higher observational coherence.

10. Categorical Descent and Observer Networks

The preceding sections have established a set-theoretic and order-theoretic account of distributed object formation. Local observational domains carry equivalence relations. Restriction transports those relations to smaller domains. Compatible local policies determine a matching family. Equivalence closure constructs the minimal global identity relation. Faithful descent occurs precisely when the generated relation restricts back to the original local policies.

Category theory reorganizes these constructions around objects, morphisms, diagrams, and universal properties.

This reorganization is not merely a change of terminology. It reveals several structures that remain partially hidden in the direct relation-based treatment.

A local observer is not only a subset equipped with an equivalence relation. It is a policy-indexed state space capable of receiving and transmitting identity-preserving maps.

A cover is not merely a list of subsets. It is an indexing diagram containing domains, overlaps, multiple intersections, and restriction morphisms.

A compatible family is not merely a collection satisfying equations. It is an element of a categorical limit called a matching object.

A faithful globalization is not merely an equivalence relation passing a restriction test. It is a realization of local descent data.

The set of all such realizations forms a category or partially ordered space of possible global ontologies.

The minimal globalization is then characterized as an initial or least realization whenever faithful descent succeeds.

This categorical viewpoint also clarifies the limits of the strict set-theoretic model. The local theories developed so far require literal equality of policies on overlaps and literal equality of local transformations where domains intersect. Many scientific and mathematical settings provide only equivalence, isomorphism, translation, or transformation between local presentations. Their local objects do not glue strictly. They glue up to specified morphisms.

Once gluing occurs up to equivalence, the descent data must record the equivalences themselves, their compatibility on triple overlaps, and eventually transformations among alternative equivalences. This leads from presheaves of sets and posets to presheaves of categories, groupoids, stacks, and higher stacks.

The present section develops the strict categorical foundation before indicating those extensions. It defines the category of policy-indexed state spaces, the restriction functors associated with domain inclusion, Čech-style diagrams of observational covers, matching objects, realization categories, morphisms between observer networks, and categorical operations of refinement and coarsening.

The central categorical structure can be summarized by a descent map

$\rho: Eq(X) \to Match(\{U_i\})$.

The codomain contains all locally compatible policy families.

The image contains the effective descent families.

The fibers contain the possible global realizations of each family.

Failure of surjectivity is gluing obstruction.

Failure of injectivity is global underdetermination.

A singleton fiber is unique globalization.

A nonempty fiber with a least element contains a canonical minimal global ontology.

Category theory places all of these possibilities inside one diagram.

10.1 Categories of Local Policy-Indexed Spaces

Define a category Obs of policy-indexed state spaces.

An object of Obs is a pair

(X, E),

where X is a set and E is an equivalence relation on X.

A morphism

f: (X, E) → (Y, F)

is a function

f: X → Y

that preserves observational identity:

x E x′ implies f(x) F f(x′).

Such a map may be called policy-compatible, identity-preserving, or observationally admissible.

The identity function on X is a morphism

idX: (X, E) → (X, E)

because x E x′ immediately implies idX(x) E idX(x′).

If

f: (X, E) → (Y, F)

and

g: (Y, F) → (Z, H)

are morphisms, then their composite

g ∘ f: (X, E) → (Z, H)

is identity-preserving.

Indeed,

x E x′

implies

f(x) F f(x′),

which implies

g(f(x)) H g(f(x′)).

Thus Obs is a category.

Every morphism in Obs induces a function between quotient object spaces:

$\bar{f}: X/E \to Y/F$,

defined by

$\bar{f}([x]E) = [f(x)]F$.

The map is well defined because equivalent representatives have F-equivalent images.

This produces a quotient functor

$Q: Obs \to Set$.

On objects,

$Q(X, E) = X/E$.

On morphisms,

$Q(f) = \bar{f}$.

The identity and composition laws are preserved:

$Q(idX) = idX/E$,

and

$Q(g \circ f) = Q(g) \circ Q(f)$.

The quotient functor formalizes the passage from policy-indexed states to effective objects.

The category Obs contains several important subcategories.

For a fixed state space X, define ObsX to be the full subcategory whose objects are pairs (X, E) with $E \in Eq(X)$.

A morphism in this full subcategory is a function $X \to X$ preserving the relevant source and target policies. If one restricts further to identity functions on the underlying state space, the category reduces to the refinement category Res(X).

The objects of Res(X) are equivalence relations on X.

There is one arrow

$E \to F$

when

$E \subseteq F$.

The arrow is induced by the identity map on X, regarded as a policy-compatible map

$id_X: (X, E) \to (X, F)$.

The condition $E \subseteq F$ ensures that every E-equivalent pair is F-equivalent.

Applying the quotient functor gives the canonical surjection

$X/E \to X/F$.

Thus refinement arrows are exactly those identity-preserving maps that leave the underlying states fixed while coarsening their object identity.

For a local domain U, the category Res(U) contains every possible local conceptual resolution.

A local observation system (U, E) is one selected object of Res(U).

A distributed observation network selects one object E_i from each local category $Res(U_i)$, subject to compatibility under restriction.

This separation is useful.

The categories $Res(U_i)$ contain all possible policies.

The observer network chooses one policy from each category.

The descent problem asks whether the selected objects arise as restrictions of one global object $E \in Res(X)$.

The category Obs can also include maps between different domains.

Suppose U and V are local state spaces with policies E and F. A map

$f: (U, E) \to (V, F)$

can represent:

a sensor transformation,

a change of representation,

a database translation,

a model comparison,

a coordinate transition,

a coarse-graining map,

or an information channel.

The identity-preservation condition ensures that states already treated as one object in U are not separated by f in V.

The map may identify additional states. Thus morphisms in Obs may reduce distinction.

If stronger preservation is required, one may demand both directions:

x E x′ if and only if f(x) F f(x′).

Such a map reflects as well as preserves identity.

If f is also bijective, it is an isomorphism of policy-indexed spaces.

Isomorphisms represent alternative presentations of exactly the same object system.

The distinction among identity-preserving maps, identity-reflecting maps, and isomorphisms becomes important in higher descent. Strict subset restrictions use literal inclusion. General observer networks may connect local contexts through noninvertible translations or invertible changes of presentation.

The category Obs is therefore the basic arena in which observation policies become transportable mathematical structures.

10.2 Restriction Functors

Let $V \subseteq U$.

Restriction assigns to every policy $E \in Eq(U)$ the policy

$E|V \in Eq(V)$.

At the level of thin resolution categories, this gives a functor

$Res(U) \to Res(V)$.

On objects,

$E \mapsto E|V$.

If $E \to F$ is a refinement arrow in Res(U), then $E \subseteq F$. Restriction preserves inclusion:

$E|V \subseteq F|V$.

Therefore there is a corresponding arrow

$E|V \to F|V$

in Res(V).

Identity arrows and composition are preserved because restriction is monotone.

Thus every inclusion of domains induces a restriction functor in the opposite direction.

For nested inclusions

$W \subseteq V \subseteq U$,

the restriction functors satisfy

$\mathrm{Res}(U) \to \mathrm{Res}(V) \to \mathrm{Res}(W)$

equal to the direct restriction

$\mathrm{Res}(U) \to \mathrm{Res}(W)$.

Restriction to the same domain is the identity functor.

The assignment

$U \mapsto \mathrm{Res}(U)$

therefore defines a contravariant functor from the category of observational domains into the category of thin categories.

Equivalently, it defines a presheaf of categories.

Because each Res(U) is a complete lattice regarded as a thin category, the assignment also carries order-theoretic structure.

Restriction preserves arbitrary meets.

It preserves the finest and coarsest policies.

It need not preserve joins.

Categorically, failure of join preservation means that restriction does not preserve every colimit present in the thin resolution categories.

A globally generated coarsening may contain identity paths that leave a subdomain. Restricting the result retains the identity consequence even though the path is no longer visible. By contrast, forming the join after restriction uses only paths available inside the smaller domain.

The restriction functor therefore preserves operations built by direct intersection but not necessarily operations built by transitive generation.

This is another expression of the local-global asymmetry.

One may also define restriction at the level of policy-indexed spaces.

Given an object (U, E) and a subset V ⊆ U, its restriction is

(V, E|V).

The inclusion map

ιV,U: V → U

is identity-preserving:

x E|V y implies x E y.

Thus there is a morphism

(V, E|V) → (U, E)

in Obs.

Notice the two directions involved.

The policy restriction functor moves from E on U to E|V on V.

The inclusion morphism in Obs moves from the smaller policy-indexed space into the larger one.

These directions are compatible because presheaf restriction is contravariant while state-space inclusion is covariant.

The distinction becomes useful in diagrams. A cover gives inclusions of intersections into local domains. Applying the resolution presheaf reverses those arrows and produces restriction maps among policy spaces.

A global policy E therefore generates a cone of restricted policies over the entire cover diagram.

10.3 Diagrams Indexed by Observational Covers

Let {Ui}i∈I be a cover of X.

The cover determines not only the individual domains but also all of their finite intersections.

For a finite nonempty set of indices

J = {i0, …, in},

define

UJ = Ui0 ∩ ⋯ ∩ Uin.

Whenever K ⊆ J, one has

UJ ⊆ UK.

Thus intersections and inclusions form a category associated with the cover.

This category may be viewed as the Čech indexing category of the observational cover.

Its low-dimensional objects are:

the domains Ui,

the pairwise overlaps Uij,

the triple overlaps Uijk,

and higher intersections.

Its morphisms are the natural inclusions from smaller intersections into larger domains.

Applying the contravariant resolution presheaf produces a diagram

Res(Ui) ⇉ Res(Uij) ⇉ Res(Uijk) ⋯

whose arrows are restriction functors.

A selected family of local policies Ei determines policies on every multiple overlap when compatibility holds.

On Uij, define

Eij = Ei|Uij = Ej|Uij.

On Uijk, define

Eijk = Ei|Uijk = Ej|Uijk = Ek|Uijk.

The same construction continues through all finite intersections.

Thus a compatible family produces a coherent assignment of policies to the entire Čech diagram.

At the set level, one may form the diagram

Eq(Ui) ⇉ Eq(Uij) ⇉ Eq(Uijk) ⋯.

A point in the product

∏i Eq(Ui)

is an arbitrary family of local policies.

The compatibility equations select those families whose images under the two restrictions to each pairwise overlap agree.

The set of compatible families can therefore be represented as an equalizer:

Match({Ui}) =
{(Ei)i ∈ ∏i Eq(Ui) :
Ei|Uij = Ej|Uij for all i, j}.

For ordinary strict equivalence relations, pairwise agreement implies agreement on all higher intersections. Thus the pairwise equalizer already captures the matching families.

The full Čech diagram remains valuable because it generalizes correctly when overlap comparisons are nontrivial morphisms rather than literal equality.

The object Match({Ui}) is the matching object of the cover.

Its elements are descent candidates.

Every global policy E ∈ Eq(X) determines an element of Match({Ui}) by restriction:

ρ(E) = (E|Ui)i.

This defines the descent map

ρ: Eq(X) → Match({Ui}).

The properties of ρ summarize the local-to-global problem.

The map is surjective exactly when every compatible local family has at least one global realization.

The map is injective exactly when global policies are uniquely determined by their local restrictions.

The map is bijective exactly when the resolution presheaf satisfies the sheaf condition on the cover.

For a particular compatible family e = (Ei)i, the fiber

$\rho^{-1}(e)$

is the set of all faithful globalizations of that family.

The fiber is empty when descent is obstructed.

It contains one element when globalization is unique.

It contains several elements when global identity is underdetermined.

This fiber description provides a compact categorical classification of distributed ontology.

10.4 Cones, Limits, and Compatible Observer Systems

A limit is a universal object representing coherent solutions to a diagram of compatibility conditions.

For the Čech diagram of local policy spaces, the matching object is a limit.

In the simplest pairwise form, consider the product of all local policy sets

$P0 = \prod_i Eq(Ui)$

and the product of all overlap policy sets

$P1 = \prod_{i,j} Eq(Uij)$.

There are two maps

$d0, d1: P0 \rightarrow P1$.

The first sends a local family to the restrictions coming from the first domain:

$d0((Ei)i)i,j = Ei|Uij$.

The second sends it to the restrictions coming from the second domain:

$d1((Ei)i)i,j = Ej|Uij$.

The matching object is the equalizer

$Match(\{Ui\}) = Eq(d0, d1)$.

It consists of exactly those families for which

$d0((Ei)) = d1((Ei))$.

Thus compatibility is a limit condition.

No global equivalence relation is required to define this limit. The matching object exists even when many of its elements fail effective descent.

A compatible observer system can therefore be understood as a cone over the overlap diagram.

Each local policy maps consistently to policies on intersections.

The cone records agreement among the observers.

A global policy E supplies a stronger cone. It maps by restriction to every local domain and every multiple overlap. All local policy data are obtained from this one apex.

The descent question asks whether a compatible local cone has a global apex in Eq(X).

Not every matching cone is induced by a global policy.

This is the categorical form of sheaf failure.

The distinction between limits and global realizations is important.

The matching object is a limit of local compatibility data.

It does not itself solve the globalization problem.

It describes all local systems that satisfy overlap equations.

The global restriction map ρ selects those matching systems that are actually realized by global equivalence relations.

One may therefore regard effective descent as a comparison between two constructions:

the global policy space Eq(X),

and the limit of the local Čech diagram Match({Ui}).

A sheaf would make the comparison an isomorphism.

The resolution presheaf generally does not.

The failure of ρ to be an isomorphism contains two independent categorical defects.

Its failure to be an epimorphism at the set level corresponds to obstructed local families.

Its failure to be a monomorphism corresponds to distinct global policies with identical local traces.

For a fixed family e ∈ Match({Ui}), the faithful descent theorem gives an explicit test for whether e lies in the image of ρ.

Construct

E* = EqCl(⋃i Ei).

Then

e ∈ Im(ρ)

if and only if

$E^* | U_i = E_i$

for every i.

Thus the abstract image problem has a concrete closure criterion.

The global realization, when it exists minimally, is not obtained by the limit alone. It is obtained by adding transitive closure and checking whether that closure preserves the limiting local data.

Category theory therefore reveals two stages:

matching is a limit problem,

effective descent is a realization problem.

The first organizes agreement.

The second certifies objecthood.

10.5 Descent Objects and Global Realizations

Fix a compatible family

$e = (E_i)_i \in \mathrm{Match}(\{U_i\})$.

Define the realization category Real(e) as follows.

Its objects are equivalence relations E on X satisfying

$E | U_i = E_i$

for every i.

Because Eq(X) is a thin category under refinement, Real(e) inherits a thin category structure.

There is an arrow

$E \to F$

when

$E \subseteq F$.

Thus Real(e) is the partially ordered family of faithful globalizations.

The category may be empty.

It is empty exactly when e fails faithful descent.

When it is nonempty, it contains the minimal globalization E^* as its least object.

Indeed, every realization E contains each Ei and therefore contains E*. If E* itself restricts faithfully, it belongs to Real(e). Hence there is a unique arrow

$E^* \to E$

for every $E \in \mathrm{Real}(e)$.

In categorical language, E* is an initial object of Real(e).

Because Real(e) is thin, the initial object is exactly the least faithful globalization.

This gives a universal characterization of the canonical global ontology.

It is the realization that makes no identity assertion beyond those forced by the local data and equivalence closure.

Other objects of Real(e) may add underdetermined global identifications.

A terminal object of Real(e), when it exists, would be the coarsest faithful globalization. It would identify every pair that can be identified without violating any local policy.

Such a terminal realization does not always exist.

Two individually faithful global coarsenings can have a join whose combined transitive consequences create a forbidden local merger. Thus the set of faithful globalizations need not be closed under joins.

By contrast, Real(e) is closed under arbitrary meets whenever it is nonempty.

Let $\{E_\alpha\}$ be faithful realizations.

Their intersection

$E = \bigcap_\alpha E_\alpha$

is an equivalence relation.

Restriction commutes with intersection:

$E|U_i = \bigcap_\alpha (E_\alpha|U_i) = \bigcap_\alpha E_i = E_i$.

Therefore E is a faithful realization.

The initial object E* can be recovered as the intersection of all faithful realizations.

The realization category therefore forms a complete meet-semilattice with a least element but need not form a complete lattice.

This structure separates forced and optional identity.

The relation E* contains all forced identifications.

An additional pair appearing in some but not all realizations is optional or underdetermined by the local system.

A pair appearing in no realization is forbidden by local fidelity or by the transitive consequences of other required distinctions.

One may define three relations on global state pairs.

Forced identity:

x and y are equivalent in every faithful realization.

Optional identity:

x and y are equivalent in some but not all faithful realizations.

Forbidden identity:

x and y are equivalent in no faithful realization.

When Real(e) is nonempty, forced identity is exactly E*.

Optional and forbidden identities depend on the structure above E*.

The realization category is therefore more informative than a single global quotient. It records the full space of globally admissible ontologies consistent with the local observers.

A unique global object system occurs when Real(e) has exactly one object.

Global underdetermination occurs when it has several.

Obstruction occurs when it has none.

This trichotomy is categorical:

empty category,

contractible thin category with one object,

or nontrivial partially ordered realization space.

The concept of a descent object can be used in two related senses.

First, E* may be called the descent object when it is a faithful initial realization.

Second, the global quotient

X/E*

may be called the effective descent object in the category of quotient object spaces.

Every local quotient Ui/Ei maps injectively into X/E* under faithful descent. The local quotient diagram therefore receives a cocone into the global quotient.

However, the global quotient is not generally a simple colimit of the local quotient sets without additional identifications and fidelity checks. A naive colimit may merge local objects through connectivity and thereby reproduce E* even when local injectivity fails.

The effective descent object must therefore be understood together with its local embedding condition.

The quotient alone records the global classes.

Faithful descent also requires the maps from local quotient objects to remain injective.

10.6 Morphisms Between Distributed Observation Networks

A distributed observation network can be represented by the data

$N = (X, \{U_i\}_{i\in I}, \{E_i\}_{i\in I})$.

To compare two networks, define a morphism that transports states, domains, and identity policies coherently.

Let

$N = (X, \{U_i\}_{i\in I}, \{E_i\})$

and

$M = (Y, \{V_a\}_{a\in A}, \{F_a\})$.

A strict network morphism from N to M consists of:

a function

$f: X \to Y$,

and an index map

$\alpha: I \to A$,

such that

$f(U_i) \subseteq V_{\alpha(i)}$

for every i,

and

x Ei y implies f(x) Fα(i) f(y)

for every x, y ∈ Ui.

The index map assigns each source observer to a target observer capable of receiving its states.

The identity-preservation condition ensures that a local object in the source is not split by the target policy.

If a state lies in several source domains, the images supplied by f are already identical because f is global. The target domain assignments must permit the resulting overlap structure to remain meaningful.

Composition is defined naturally.

Suppose

(f, α): N → M

and

(g, β): M → L.

Then

(g ∘ f, β ∘ α): N → L

is a network morphism.

Domain containment follows:

f(Ui) ⊆ Vα(i),

and

g(Vα(i)) ⊆ Wβ(α(i)).

Identity preservation follows by composition of policy-compatible maps.

Identity morphisms use the identity function on X and the identity map on the index set.

Thus distributed observation networks form a category under strict network morphisms.

Every network has a generated global equivalence

E*N = EqCl(⋃i Ei).

A network morphism preserves each generating relation into a target local relation and therefore into the target generated relation.

If x E*N y, a zig-zag of source local identities connects x and y. Applying f produces a zig-zag of target local identities connecting f(x) and f(y).

Hence

f(x) E*M f(y).

Therefore f induces a global quotient map

$\bar{f}$: X/E*N* → *Y/E*M.

The assignment

N ↦ X/E*N

is therefore functorial on the category of strict observer networks.

This is the global object functor.

It converts a network of local observers into its minimally generated global object space.

If one restricts to faithfully descending networks, the functor maps networks to global quotient spaces whose local object maps are injective.

A network morphism can preserve global objects without preserving every local distinction. The definition requires source equivalence to map into target equivalence, but not the converse.

A stronger morphism may be required when identity reflection matters:

f(x) Fα(i) f(y) implies x Ei y.

Such a morphism embeds the source local quotient faithfully into the target local quotient.

If f is bijective on states, α is appropriately invertible, and all local identities are preserved and reflected, the networks are isomorphic.

Network isomorphism represents a change of distributed presentation without change of the underlying object structure.

More general network morphisms may require:

one source domain mapping to several target domains,

several source observers merging into one target observer,

relations rather than functions between state spaces,

or functors between categories of local states.

These extensions are natural but not required for the strict foundation.

The important result is that generated globalization is compatible with identity-preserving network maps.

Distributed object formation is therefore functorial.

A lawful transformation of observer networks induces a lawful transformation of their global object spaces.

10.7 Refinement and Coarsening Across Networks

Two observer networks on the same state space and cover may differ only in their local policies.

Let

$N = (X, \{U_i\}, \{E_i\})$

and

$N' = (X, \{U_i\}, \{F_i\})$.

Suppose

$E_i \subseteq F_i$

for every i.

Then N is pointwise finer than N′.

Every identity accepted by N is accepted by N′, while N preserves at least as many local distinctions.

The generated global relations satisfy

$E^* \subseteq F^*$.

This follows from monotonicity of union and equivalence closure.

Thus pointwise coarsening of local policies produces global coarsening.

There is a canonical quotient map

$X/E^* \to X/F^*$.

At the network level, pointwise refinement can be represented by a natural transformation between the selected local-policy diagrams.

Because the resolution categories are thin, the components are the refinement arrows

$E_i \to F_i$.

Naturality expresses compatibility with restriction.

If both families are overlap-compatible, the refinement arrows commute with all restriction maps.

The pointwise order on local policy families therefore induces an order on observer networks with fixed cover.

Several important asymmetries occur.

Coarsening can create new descent obstructions.

A family $\{E_i\}$ may descend faithfully, while a pointwise coarsening $\{F_i\}$ creates new cross-domain paths that return to a local domain and merge F_i-distinct classes.

Coarsening each local policy does not guarantee that the family remains globally faithful.

Refinement can remove obstructions by deleting identity edges.

However, independently chosen refinements may destroy overlap compatibility. If two observers refine their common overlap differently, direct disagreement appears.

A repair by refinement must therefore be coordinated as a natural transformation of the entire local diagram.

One can distinguish four operations.

Local policy refinement changes E_i while keeping U_i fixed.

Local policy coarsening changes E_i in the opposite direction.

Cover refinement replaces domains with smaller subdomains.

Cover coarsening replaces them with larger domains or adds domains containing more states.

These operations affect descent differently.

Policy refinement removes direct identity edges and can split global components.

Policy coarsening adds edges and can fuse components.

Cover refinement may hide distinctions by ensuring that certain state pairs no longer occur together in any domain.

Cover coarsening increases direct comparison and may expose or resolve hidden identity questions.

Categorically, cover refinement induces a functor between the corresponding Čech indexing categories. Policies can be pulled back to the finer cover by restriction.

If a global policy exists, its restrictions to any cover refinement remain compatible.

Conversely, descent on a very fine cover need not imply descent on a coarser cover when the coarser local policies contain stronger identity or distinction requirements.

The choice of cover therefore participates in the ontology.

A policy family cannot be evaluated independently of the observational architecture that supports it.

One may also define a category of covers of X ordered by refinement. The resolution presheaf and its matching objects then vary over this cover category.

A change of cover induces comparison maps between local descent problems.

This larger structure resembles a fibred category.

Over each cover lies a category of compatible policy families.

Over each compatible family lies a realization category.

Moving between covers transports local policies and changes the constraints on global identity.

Such a construction would support a systematic theory of observer-network design.

One could ask:

Which cover detects a given gluing obstruction?

Which additional domain makes globalization unique?

Which cover minimizes observational cost while preserving faithful descent?

Which refinements isolate an erroneous bridge?

Which coarsenings supply enough pair comparisons to eliminate underdetermination?

These are categorical optimization problems over the architecture of observation itself.

10.8 Toward Stacks and Higher Observational Coherence

The strict theory assumes that local domains are literal subsets of one global set X and that shared states are literally equal on overlaps.

It also assumes that local policies agree by equality:

Ei|Uij = Ej|Uij.

Many important systems do not satisfy these assumptions.

Two laboratories may not share literal samples but may possess records connected by a calibration or identification map.

Two coordinate charts use different representations of the same geometric region.

Two scientific theories describe a common phenomenon through different vocabularies.

Two databases contain separate tokens intended to refer to the same entity.

Two AI models possess latent representations that correspond only through learned transformations.

In these settings, local object systems may agree up to isomorphism rather than equality.

Suppose each domain Ui carries a category Ci of local objects rather than merely a quotient set.

On an overlap Uij, there may be an equivalence of categories

Φij: Cj|Uij → Ci|Uij.

The reverse comparison may be supplied by a quasi-inverse.

On triple overlaps, the composites

Φij ∘ Φjk

and

Φik

may not be literally equal. They may be connected by a specified natural isomorphism.

On quadruple overlaps, those natural isomorphisms must themselves satisfy coherence laws.

This is the beginning of stack theory.

A stack is a categorified sheaf in which local objects and morphisms glue up to coherent isomorphism.

The transition from the present theory to stacks is motivated by several limitations of set-theoretic quotienting.

First, quotienting erases provenance.

Once states enter one equivalence class, the quotient does not record which local observations established their identity.

A groupoid can retain states as objects and identity witnesses as invertible morphisms.

Second, quotienting erases multiplicity of paths.

Two states may be connected by several independent chains of evidence. The set quotient records only that they belong to one class.

A path groupoid can retain distinct equivalence paths.

Third, quotienting makes every established equivalence indistinguishable from literal equality at the quotient level.

In many contexts, it matters whether two presentations are equal, isomorphic, empirically matched, causally linked, or merely similar.

An enriched category can label morphisms by type, confidence, cost, or provenance.

Fourth, strict overlap equality may be too rigid.

Local classifications may differ in labels or presentation while encoding equivalent object structures.

Stack-like descent allows those presentations to glue through specified isomorphisms.

Fifth, false gluing can result from collapsing nontrivial equivalence structure too early.

Two local objects may be related by a reversible transformation without needing to become one literal quotient point.

A groupoid retains both objects and the equivalence between them.

The set of connected components can still be formed later if appropriate.

Sixth, contextual systems may lack one global assignment even though all local data and transformations form a coherent higher object.

Quantum contextuality provides a broad motivating pattern. Local measurement contexts may possess internally consistent outcome structures without admitting one global assignment satisfying all local constraints.

A higher observational object may encode the contextual system even when a global set-valued section does not exist.

Seventh, alternative gluing procedures may themselves be related.

One pair of local objects may admit several comparison maps. Homotopies or higher morphisms can express when those comparisons are equivalent.

The hierarchy of possible observational structures is therefore:

sets of quotient objects,

groupoids of presentations and equivalences,

categories of possibly noninvertible transformations,

stacks of locally defined categories,

2-stacks with morphisms among descent equivalences,

and higher stacks retaining coherence at all levels.

The present equivalence-relation theory occupies the first strict layer.

Its obstruction theorem identifies exactly when connected-component collapse fails to preserve local identity.

A groupoid extension would retain the local states and identity paths before collapse.

Let G be the graph generated by local equivalence edges.

Instead of immediately replacing each connected component by one quotient point, form the path groupoid Π(G).

The objects are states.

Morphisms are paths of local identities, modulo a chosen path equivalence.

Composition is concatenation.

Inverses are path reversal.

The ordinary global quotient X/E* is then the set of connected components of this groupoid.

Thus set-theoretic globalization is a decategorification.

It forgets the morphisms and retains only component membership.

False gluing may indicate that this decategorification was premature.

The correct global object may be the groupoid itself rather than its component set.

Similarly, when local quotient categories are connected through equivalences on overlaps, the correct descent object may be a stack rather than one global quotient category.

A stack does not force all equivalent local presentations to become literally identical. It retains their equivalences as structure.

This creates a more discriminating account of objecthood.

An object can possess several locally valid presentations.

Those presentations can be related without being collapsed.

Alternative paths of comparison can be recorded.

Context can remain visible.

Global coherence can exist at the level of equivalence even when strict global identity fails.

The move toward stacks also changes the interpretation of gluing obstruction.

In the set-based theory, an obstruction means no global equivalence relation reproduces the local policies exactly.

In a stack-like theory, strict descent may fail while weak descent succeeds.

The local objects may glue up to coherent equivalence even though they do not glue into one literal global set partition.

Thus some apparent failures of object formation may be failures only of an overly strict ontology.

The correct response may not be to revise local data or reject global coherence.

It may be to raise the categorical level.

The resulting research program proceeds through several stages.

First, replace quotient sets by groupoids of states and identity witnesses.

Second, define restriction functors between local groupoids.

Third, express overlap agreement through equivalences rather than equality.

Fourth, impose cocycle coherence through natural isomorphisms on triple overlaps.

Fifth, define effective descent objects in a stack of observational groupoids.

Sixth, introduce noninvertible morphisms for abstraction, refinement, and transformation.

Seventh, enrich identity paths with confidence, provenance, cost, or admissibility type.

Eighth, develop higher morphisms among alternative observational and explanatory paths.

These extensions are not decorative generalizations. They are demanded by the phenomena already found in the strict theory.

Pairwise compatibility can fail to imply faithful global identity.

Equivalence closure can erase locally required distinctions.

Several gluing paths can support different interpretations.

Local presentations can be equivalent without being identical.

The global object may exist only as a coherent network of relations.

Category theory therefore changes the final question.

The set-theoretic question is:

Which states belong to the same global equivalence class?

The categorical question is:

What system of local objects, transformations, equivalences, and higher coherences constitutes the global object?

The first question produces quotient ontology.

The second produces relational ontology.

Both are necessary.

The quotient is appropriate when internal paths and presentations can be forgotten safely.

The higher object is required when the structure of identification remains scientifically or mathematically significant.

The categorical results of this section can now be summarized.

Policy-indexed state spaces form a category Obs.

Quotient object formation defines a functor from Obs to sets.

Every domain U carries a thin category Res(U) of conceptual resolutions.

Domain inclusions induce contravariant restriction functors.

An observational cover determines a Čech-style diagram of local resolution categories.

The limit of the local restriction diagram is the matching object of compatible policy families.

Global policies map into that matching object through restriction.

The fibers of this descent map are realization categories.

An empty fiber represents obstruction.

A singleton fiber represents unique globalization.

A nontrivial fiber represents underdetermination.

When the fiber is nonempty, the generated global equivalence E* is its initial object.

Morphisms between observer networks induce maps between generated global quotient spaces.

Refinement and coarsening become natural transformations among local-policy diagrams.

Strict set-valued descent is only the first layer.

When local objects agree only up to equivalence, stacks and higher categories provide the appropriate extension.

The next section applies the full theory to explicit finite and scientific constructions. It develops examples of faithful descent, obstruction, false fusion, distributed classification, multisensor tracking, biological integration, multiscale physics, and algorithms for detecting gluing defects.

11. Examples and Computational Constructions

The local-to-global theory becomes most transparent when the abstract definitions are applied to explicit systems. Finite examples reveal the exact role of overlap agreement, transitive closure, returning zig-zags, and local object fusion. Applied examples show how the same structures arise in image recognition, sensor networks, biological classification, scientific integration, and multiscale physical modeling.

The examples also clarify an important distinction between mathematical diagnosis and domain judgment.

The mathematics can determine:

which local identities generate a common global class,

whether the resulting globalization is faithful,

which local objects are fused,

which paths cause the fusion,

and which identity declarations must be revised to restore descent.

It cannot determine from equivalence structure alone whether a generated fusion is scientifically correct. That decision requires information about causation, dynamics, intervention, measurement reliability, or task-specific admissibility.

The computational problem is nevertheless exact once the local policies are specified. For finite systems, minimal globalization reduces to connected-component computation. Faithful descent reduces to checking whether one connected component contains two distinct local equivalence classes from the same domain. Every failure can be accompanied by a finite witness path.

This section begins with a small system in which distributed identity glues faithfully. It then presents a compatible family that cannot descend, followed by a false-fusion construction caused by one bridge. The later subsections translate the formal structure into practical observer networks and conclude with algorithms for detecting, explaining, and repairing gluing defects.

11.1 A Three-Domain Finite Gluing Example

Let the global state space be

X = {a, b, c, d, e}.

Define three observational domains:

U1 = {a, b, c},

U2 = {b, c, d},

U3 = {a, d, e}.

Assign the following local equivalence classes.

On U1, let E1 have classes

{a, b},

{c}.

Observer 1 identifies a and b while distinguishing c.

On U2, let E2 have classes

{b},

{c, d}.

Observer 2 identifies c and d while distinguishing b.

On U3, let E3 have classes

{a},

{d, e}.

Observer 3 identifies d and e while distinguishing a.

The pairwise overlaps are

$U1 \cap U2 = \{b, c\}$,

$U1 \cap U3 = \{a\}$,

$U2 \cap U3 = \{d\}$.

On the overlap {b, c}, both E1 and E2 distinguish b from c.

On the singleton overlaps {a} and {d}, compatibility is automatic.

Thus the family is pairwise compatible.

The direct local identity edges are

a E1 b,

c E2 d,

d E3 e.

The generated global equivalence relation E* has two nonempty classes:

$C1 = \{a, b\}$,

$C2 = \{c, d, e\}$.

There is no local identity path connecting the two components.

Now restrict E* to each domain.

On U1,

$C1 \cap U1 = \{a, b\}$,

$C2 \cap U1 = \{c\}$.

These are exactly the E1-classes.

On U2,

$C1 \cap U2 = \{b\}$,

$C2 \cap U2 = \{c, d\}$.

These are exactly the E2-classes.

On U3,

$C1 \cap U3 = \{a\}$,

$C2 \cap U3 = \{d, e\}$.

These are exactly the E3-classes.

Therefore

$E^*|U1 = E1$,

$E^*|U2 = E2$,

$E^*|U3 = E3$.

Faithful descent holds.

The local quotient spaces are

$U1/E1 = \{\{a, b\}, \{c\}\}$,

$U2/E2 = \{\{b\}, \{c, d\}\}$,

$U3/E3 = \{\{a\}, \{d, e\}\}$.

The global quotient is

$X/E^* = \{\{a, b\}, \{c, d, e\}\}$.

Each local quotient maps injectively into the global quotient.

For U1, the local class {a, b} maps to C1, while {c} maps to C2.

For U2, {b} maps to C1, while {c, d} maps to C2.

For U3, {a} maps to C1, while {d, e} maps to C2.

The maps are not injective across the disjoint union of all local quotient spaces. Several local classes from different domains map to the same global object. That is the intended gluing.

The class {a, b} from U1, the singleton {b} from U2, and the singleton {a} from U3 are three local traces of C1.

The singleton {c} from U1, the class {c, d} from U2, and the class {d, e} from U3 are traces of C2.

No local domain contains all of C2. The global object {c, d, e} exists only through the chain

c E2 d E3 e.

Observer 2 sees c and d as one object.

Observer 3 sees d and e as one object.

Their shared state d transmits identity.

Observer 1 sees only c from that global object.

The class C2 is therefore genuinely distributed but faithfully formed.

This example illustrates several general principles.

A global object need not be represented completely by any one observer.

Different observers may see differently sized traces of the same global object.

Faithful descent permits fusion of local objects from different domains.

It prohibits fusion of distinct local objects within one domain.

Singleton overlaps can transmit identity even though they contain no nontrivial pairwise relation to compare.

The final global object system can be strictly smaller than the total collection of local objects without contradicting any local policy.

The example also shows why the global quotient should not be constructed by treating all local quotient classes as permanently distinct. Local objects are presentations. Globalization determines which presentations refer to the same distributed object.

11.2 A Compatible Family That Fails Faithful Descent

Let

X = {a, b, c, d}.

Define

U1 = {a, b, d},

U2 = {b, c},

U3 = {c, d}.

Let E1 have classes

{a, b},

{d}.

Let E2 have one class

{b, c}.

Let E3 have one class

{c, d}.

The pairwise overlaps are

U1 ∩ U2 = {b},

U2 ∩ U3 = {c},

U1 ∩ U3 = {d}.

Every overlap is a singleton. All local policies therefore agree on all pairwise and higher intersections.

The family is compatible.

The direct local identity edges are

a E1 b,

b E2 c,

c E3 d.

The generated relation contains the zig-zag

a E1 b E2 c E3 d.

Therefore

a E* d.

Indeed, every state belongs to one generated class:

X/E* = {{a, b, c, d}}.

Restrict this global relation to U1.

The restriction identifies all three states a, b, and d.

But E1 has two classes:

{a, b},

{d}.

Thus

$E1 \subset E^*|U1$.

Faithful descent fails.

The defect relation O1 contains every ordered pair crossing the two original E1-classes:

(a, d),

(d, a),

(b, d),

(d, b).

The local object-loss measure is

$L1 = |U1/E1| - |U1/(E^*|U1)|$

$= 2 - 1$

$= 1$.

One local object distinction has been destroyed.

The shortest obstruction path from b to d is

b E2 c E3 d.

It has length two and traverses two external local policies.

The shortest path from a to d is

a E1 b E2 c E3 d.

It has length three.

The failure is invisible to pairwise overlap testing.

No observer contradicts another on a shared pair.

The problem appears only when the three locally valid relations are composed.

No alternative global equivalence relation can repair the system while retaining all local identities. Any global relation containing E1, E2, and E3 must identify a, b, c, and d.

The obstruction can be repaired in several ways.

Remove the identity b E2 c.

Then the generated classes become

{a, b},

{c, d}.

All local restrictions are faithful if E2 is correspondingly refined to distinguish b and c.

Remove c E3 d.

Then the generated classes become

{a, b, c},

{d}.

Again E3 must be refined.

Refine E1 no further; E1 is not the source of the connecting path, although it supplies one endpoint class.

Coarsen E1 by merging d with {a, b}.

Then the original global closure becomes faithful relative to the revised policy. This repair accepts the generated identity rather than breaking it.

Add a new observational domain containing b and d with a policy that distinguishes them.

This does not by itself alter the existing positive identity chain. It makes the conflict explicit. A policy-selection procedure must then decide which identity edge to reject.

Represent one of the shared states contextually rather than literally. If c in U2 and c in U3 are not the same state but two context-specific presentations connected by a nonidentity comparison, the automatic bridge may disappear.

Each repair encodes a different interpretation.

Edge deletion says that one local equivalence was invalid.

Local coarsening says that the global system revealed a broader identity.

Domain duplication says that a shared state was contextually ambiguous.

The obstruction theorem identifies where a decision is required but does not choose the decision.

11.3 False Fusion Produced by a Transitive Zig-Zag

A single bridge can merge two large and otherwise coherent object systems.

Let X contain two groups of states:

A = {a1, a2, a3, a4},

B = {b1, b2, b3, b4}.

Suppose several local policies correctly establish that all states in A are presentations of one object and all states in B are presentations of another object.

The generated relation before bridging has two classes:

A,

B.

Now introduce a local domain

V = {a4, z, b1}

containing an intermediate state z.

Let the local policy on V identify

a4 with z

and

z with b1.

Because the policy is an equivalence relation on V, it must also identify a4 with b1. Alternatively, the two edges may be supplied by two overlapping domains. In either case, the global identity graph acquires a path from A to B.

Equivalence closure merges the two components:

A ∪ B ∪ {z}.

The resulting fusion may be mathematically unavoidable but scientifically false.

Suppose another domain W contains a2 and b3 and explicitly distinguishes them as two simultaneously existing people, two different physical objects, two unrelated records, or two biologically separate specimens.

Then the returning global path creates a descent defect on W.

The bridge has a disproportionate effect.

One local error involving a4, z, and b1 does not merely misclassify those three states. It collapses every presentation in A and every presentation in B.

This amplification is an inherent property of equivalence closure.

Identity is component-wide.

Once one edge joins two components, every member of one becomes equivalent to every member of the other.

Let |A| = m and |B| = n.

Before fusion, there are two global objects.

After fusion, there is one.

At the pair level, the bridge introduces 2mn new ordered cross-component equivalences, excluding any additional pairs involving z.

Thus one accepted bridge can generate a quadratic number of new pairwise identities.

This is why large identity systems are vulnerable to bridge errors.

The appropriate diagnostic object is not only the final defect-pair count. It is the bridge sensitivity of each accepted identity edge.

Remove the edge or local relation involving z and recompute the connected components.

If A and B separate, the edge is a cut bridge in the identity graph.

Its removal restores the previous ontology.

In a weighted system, one may assign confidence values to identity declarations. The most weakly supported edge on a false-fusion path becomes a candidate for rejection.

However, selecting the lowest-confidence edge is not always sufficient.

Several redundant paths may connect A and B.

The fusion may persist after one edge is removed.

A minimum repair set is a collection of identity edges whose removal separates every locally forbidden pair.

This is a constrained cut problem.

The false-fusion example also demonstrates the importance of negative information.

The positive identity graph contains evidence for sameness.

A locally certified distinction between a2 and b3 acts as a negative constraint.

Without explicit negative constraints, the system cannot distinguish false fusion from acceptable global merger. It will simply compute connected components.

A practical identity system should therefore retain both:

positive equivalence declarations,

and authoritative separation declarations.

Faithful descent is equivalent to consistency between their transitive consequences.

11.4 Distributed Image Classification

Consider a distributed image system in which several classifiers operate on different views, sensors, datasets, or feature spaces.

One model processes frontal images.

A second processes profile images.

A third processes low-light images.

A fourth processes thermal images.

Each model forms local equivalence classes among its inputs. Images assigned to one tracked or recognized object belong to the same local class.

The local domains may overlap through shared images, synchronized captures, common identifiers, or paired observations.

Let Ui be the images processed by classifier i.

Let Ei identify images treated as presentations of the same object.

The local quotient

Ui/Ei

is the set of object identities recognized by that classifier.

A global recognition system attempts to determine which local classes across models correspond to the same real object.

Shared observations generate links.

If a frontal image and a thermal image are paired at acquisition, their local classes may be connected.

If a profile image and a low-light image share a trusted event identifier, their classes may be connected.

The global identity graph has local object classes as nodes and cross-model correspondences as edges.

Connected components become candidate global objects.

Faithful descent requires that one connected component contain at most one class from each local classifier.

Suppose classifier 1 distinguishes two people as classes A and B.

Classifier 2 contains a low-quality image class C matched to A.

Classifier 3 matches C to a class D.

Classifier 4 matches D to B.

The chain

A — C — D — B

merges A and B globally.

Because A and B are distinct classes in classifier 1, the component violates faithful descent.

The classification system has created false gluing unless the global evidence justifies revising classifier 1.

This framework exposes a limitation of pairwise matching benchmarks.

Every cross-model match may appear acceptable when evaluated independently.

The integrated component may still be impossible.

Evaluation must therefore include component-level consistency.

Several practical constraints can be incorporated.

A cannot-link constraint states that two local classes must remain globally distinct.

A must-link constraint states that two classes must belong to one global object.

A confidence score attaches strength to a proposed correspondence.

A type constraint states that only certain classes of identity evidence may compose.

A temporal constraint prevents one object from occupying incompatible locations at the same time.

A geometric constraint tests whether views are physically compatible.

A causal constraint requires that linked images can arise from one continuous object history.

The exact equivalence framework handles must-link and cannot-link constraints cleanly.

Weighted and typed extensions are needed for uncertain evidence.

The local-to-global theory also suggests a better output than a single forced identity label.

For each candidate global object, the system can retain:

its local class traces,

the paths connecting them,

the confidence and provenance of each edge,

the local distinctions threatened by the component,

and the transformations under which the identity remains stable.

This representation supports auditability.

A reviewer can determine why two images were fused and which local classifier disagreed.

The system can distinguish:

a high-confidence global identity supported by independent paths,

an underdetermined relation between disconnected components,

and an obstructed component containing a known local contradiction.

Distributed image classification therefore becomes a descent problem rather than merely a nearest-neighbor aggregation problem.

11.5 Multisensor Object Tracking

A multisensor tracking system observes moving objects through spatially and temporally restricted fields of view.

Each sensor produces local detections and local tracks.

A local track is already a quotient object. It identifies several detections across time as presentations of one persisting object.

Let Ui contain detections accessible to sensor i during a selected interval.

Let Ei identify detections assigned to the same local track.

The quotient Ui/Ei is the set of local tracks.

Overlapping sensors create shared or comparable detections. Handoff algorithms propose that a track ending in one sensor corresponds to a track beginning in another.

These handoffs create edges among local track objects.

A connected component in the handoff graph becomes a global trajectory.

Faithful descent requires that no global trajectory contain two local tracks that the same sensor distinguishes as different objects when those tracks are intended to coexist separately.

Suppose sensor A observes two vehicles, producing local tracks A1 and A2.

Sensor B receives a handoff from A1 and labels the continuation B1.

Sensor C links B1 to C1.

A later handoff connects C1 to A2 when the objects return to sensor A’s region.

The global chain is

A1 — B1 — C1 — A2.

The resulting component contains two different A-tracks.

If A1 and A2 represent simultaneously distinct vehicles, the global trajectory is impossible.

The returning path is an obstruction witness.

This case illustrates why loop closure in tracking is both valuable and dangerous.

A correct loop can recognize that a moving object has returned.

An incorrect loop can merge distinct objects.

Local descent constraints provide a precise test.

The system should also impose dynamical stability.

A proposed global trajectory must support a coherent motion model. Position, velocity, timing, and physical reachability should be compatible along the identity path.

An identity component that is topologically connected but dynamically impossible should be rejected.

Thus tracking requires two levels of validation:

static identity descent,

and transformation consistency under motion.

One may represent time evolution by local maps or relations between detection sets.

A global track class is stable when different local representatives lead to compatible future track classes.

If two detections assigned to one global object imply mutually exclusive future positions, representative-independent dynamics fail.

The quotient object is then unstable.

Multisensor tracking also demonstrates why exact equivalence may need temporal indexing.

An object at time t and the same object at time t + 1 are not identical physical states. They are linked through persistence.

The equivalence relation represents identity across state variation, while the transformation represents temporal evolution.

A richer model may use a groupoid of observations and track morphisms rather than collapsing every detection immediately into one class.

The final connected components can still represent global object identities, but the path structure remains available for trajectory reconstruction and error analysis.

A practical tracking pipeline derived from the theory would:

form local tracks,

retain explicit must-separate relations among simultaneous local tracks,

construct candidate cross-sensor handoffs,

compute connected components,

detect any component containing forbidden local track pairs,

test dynamical feasibility,

remove or revise low-admissibility handoffs,

and certify the surviving global trajectories.

This is a direct computational realization of descent plus stability.

11.6 Biological Taxonomies Assembled Across Laboratories

Biological classification is often distributed across laboratories, geographic regions, measurement platforms, and levels of organization.

One laboratory may classify specimens by morphology.

Another may use genomic sequence.

Another may use ecological role.

Another may use reproductive compatibility.

Another may use developmental history or biochemical phenotype.

Each local system determines an equivalence relation relative to its chosen criteria.

Let Ui be the specimens or records available to laboratory i.

Let Ei identify specimens assigned to the same local biological class.

The class may represent a strain, lineage, species candidate, cell type, disease subtype, ecological guild, or another effective biological object.

Shared specimens, reference samples, sequence records, or identifiers create overlaps.

The local policies may agree on all shared cases while still generating a global obstruction.

Consider a hypothetical geographic chain of populations

P1, P2, P3, P4.

Laboratory 1 studies P1 and P2 and groups them together.

Laboratory 2 studies P2 and P3 and groups them together.

Laboratory 3 studies P3 and P4 and groups them together.

A fourth laboratory studies P1 and P4 under a criterion that distinguishes them.

The positive chain forces

P1 E* P4.

The fourth policy forbids that identification.

The result is not resolved by declaring transitivity mathematically invalid. The equivalence closure is correct relative to the local sameness declarations.

The biological issue is whether all declarations represent the same kind of identity.

Neighboring reproductive compatibility, morphological similarity, shared ancestry, and ecological interchangeability may not compose into one universal equivalence relation.

The obstruction reveals that biological identity is policy-indexed.

Different purposes may require different quotients.

A lineage policy may merge the populations.

An ecological policy may distinguish them.

A reproductive policy may fail to be globally transitive when represented through local pairwise compatibility.

In such a case, forcing the biological relation into one equivalence relation may itself be the modeling error.

The appropriate structure could be:

a graph of local compatibility,

a category of transformations,

a genealogy or phylogeny,

a graded similarity space,

or several policy-indexed equivalence relations.

The descent defect identifies where the equivalence model becomes inadequate.

A second biological example concerns cell-type integration across laboratories.

Each laboratory clusters cells measured by a different platform. Shared reference samples or marker profiles connect clusters across datasets.

The global integration graph may reveal a cell type not captured completely in any one experiment.

This can be emergent global identity.

However, one weak marker correspondence can also merge two functionally distinct cell types.

Transformation and intervention provide certification.

If the proposed global class responds coherently to perturbation, follows one developmental trajectory, or supports consistent predictive behavior, the merger gains validity.

If its members respond differently in ways essential to the study, the fusion is destructive.

The theory therefore separates three tasks that are often conflated:

forming a connected integrated cluster,

testing whether the cluster descends faithfully to local classifications,

and testing whether it behaves as one biological object.

The first is graph closure.

The second is descent.

The third is stability and scientific certification.

A rigorous biological integration system should report all three.

11.7 Multiscale Physical Descriptions

Multiscale physics involves several levels of description, each with its own state variables, equivalence policies, and admissible dynamics.

A microscopic model may distinguish detailed molecular configurations.

A mesoscopic model may identify configurations producing the same collective mode.

A macroscopic model may identify all microstates with the same thermodynamic variables.

These levels are not always literal subsets of one common state space. They may be connected by coarse-graining maps.

Let Xmicro be a microscopic state space.

Let Xmeso be a mesoscopic state space.

Let Xmacro be a macroscopic state space.

A coarse-graining map

c1: Xmicro → Xmeso

groups microscopic states according to mesoscopic observables.

A second map

c2: Xmeso → Xmacro

groups mesoscopic states according to macroscopic variables.

Each map induces a kernel equivalence relation.

On Xmicro,

$x \, E_1 \, y$

when

$c_1(x) = c_1(y)$.

The composite

$c_2 \circ c_1$

induces a coarser microscopic relation E_2.

One has

$E_1 \subseteq E_2$.

This is a vertical hierarchy of resolution rather than a cover by subsets.

Nevertheless, local-to-global issues arise when several coarse-graining schemes overlap.

Suppose one model groups microstates by energy and particle number.

Another groups them by spatial density fields.

A third groups them by transport behavior.

Each policy is valid for a restricted family of questions.

Attempting to combine all equivalences into one global relation may generate destructive fusion.

A state equivalent to another under energy variables may be equivalent to a third under density variables. Transitive closure can identify the first and third even when they differ in transport behavior that one model preserves.

The problem is formally analogous to distributed gluing, although the domains are observational contexts rather than subsets.

The generalized structure consists of several maps or functors from a common state system into local model spaces.

Their kernel relations generate candidate global identity.

Faithfulness asks whether the generated relation preserves every distinction required by each model.

This can fail.

Multiscale consistency therefore requires more than agreement of numerical predictions on overlap regimes. It requires compatibility of identity policies.

A coarse variable can be dynamically valid for one time scale but unstable for another.

Two microstates assigned to one macrostate may later evolve into different macrostates.

The quotient then fails transformation invariance.

A multiscale object is lawful when:

its local definitions agree where models overlap,

its global identity does not erase distinctions required by any contributing scale,

and its quotient remains stable under the effective dynamics.

Phase transitions provide an especially important setting.

A coarse policy appropriate in one regime may cease to be stable near a critical boundary. New distinctions become dynamically relevant. The equivalence relation must refine.

Conversely, after relaxation, many detailed states may become effectively interchangeable, allowing coarsening.

The object ontology therefore changes with scale, parameters, and dynamical regime.

The local-to-global framework can represent this through a family of policy-indexed state spaces and morphisms among them.

A full treatment may require categories rather than one set cover.

The strict theory nevertheless supplies the central diagnostic:

When several effective descriptions are combined, test whether the equivalence closure of their licensed substitutions returns to any model and erases a distinction that model requires.

If it does, the models do not inhabit one unqualified quotient ontology.

Their identities must remain indexed by scale or context.

11.8 Algorithms for Detecting Gluing Defects

For a finite distributed observation system, faithful descent can be decided efficiently.

The input consists of:

a finite state set X,

a finite collection of domains Ui,

and a partition of each Ui representing Ei.

The partitions may be given as class labels, lists of equivalence classes, or local identity edges.

The basic algorithm has four stages.

First, encode all local equivalence classes.

Second, compute the generated global components.

Third, compare global components with local classes.

Fourth, recover witnesses for any detected collision.

The most direct implementation uses a disjoint-set union structure, also called union-find.

Initialize one global component for every state $x \in X$.

For each local equivalence class A in every Ei, union all states in A.

After all local classes have been processed, two states belong to the same disjoint-set component exactly when they are E*-equivalent.

This computes the minimal globalization.

The next step tests faithful descent.

For every domain Ui, assign each local Ei-class a class identifier.

For every state $x \in Ui$, compute:

its local class identifier li(x),

and its generated global component identifier gi(x).

Create a map from each pair

(i, gi(x))

to the local class identifiers encountered inside that global component.

Faithful descent holds exactly when every such set of local identifiers has size at most one.

Equivalently, for each domain i and each global component C, all states of $C \cap Ui$ must share one local class label.

If two different local class labels occur, the component produces a gluing defect.

This test avoids enumerating all pairs in Oi.

Its complexity is approximately linear in the total size of the local class representation, up to the inverse-Ackermann factor associated with union-find operations.

Let

$N = |X|$

and let

$M = \Sigma i \ |Ui|$

count state occurrences across domains.

If each local class is unioned through one representative, the globalization stage requires $O(M \alpha(N))$ time, where α is the inverse-Ackermann function.

The fidelity test also requires $O(M \alpha(N))$ time.

Thus finite faithful descent is computationally tractable even for large systems, provided the local equivalence classes are explicitly available.

A second implementation uses the local-object graph.

Create one node for every local equivalence class in every domain.

For each state x appearing in several domains, connect all local-class nodes containing x.

Compute connected components of this graph.

A descent defect occurs when one component contains two distinct nodes from the same domain.

This representation may be smaller than the state graph when local classes contain many states.

It also aligns directly with applications in which local objects, rather than individual states, are the primary data.

However, care is needed when one state occurs in several local classes only because the policies are incompatible on overlaps. Direct overlap conflict should be detected before or during graph construction.

A complete diagnostic algorithm can be organized as follows.

Step 1: Validate local equivalence.

Confirm that each supplied local policy is a partition of Ui or otherwise satisfies reflexivity, symmetry, and transitivity.

Step 2: Check overlap compatibility.

For every pair of domains with nonempty intersection, verify that shared states receive the same local partition relation.

If a direct conflict exists, report failed gluing before proceeding.

Step 3: Construct minimal globalization.

Use union-find on state-level local classes or connected components on the local-object graph.

Step 4: Test local injectivity.

For every domain, confirm that no generated component contains more than one local class.

Step 5: Record defect classes.

For each violation, report:

the affected domain,

the local classes fused,

the generated global component,

and the states involved.

Step 6: Recover obstruction witnesses.

The union-find structure alone reports connectivity but not necessarily a path. Maintain a provenance graph containing the local identity edges used to create unions.

For two collided local classes, run a shortest-path search between selected representatives inside the provenance graph.

The resulting path is a zig-zag certificate.

Step 7: Compress the witness.

Remove redundant steps occurring inside one local equivalence class. Report the sequence of local object classes and domains traversed.

Step 8: Identify repair candidates.

Compute identity edges whose removal separates the forbidden local classes.

Exact minimum repair may be computationally difficult in large general systems, but useful approximations include:

removing the lowest-confidence edge on a witness,

finding graph bridges,

computing minimum cuts between forbidden class pairs,

or solving a weighted constraint optimization problem.

Step 9: Recompute closure after proposed repair.

A local edge may participate in several obstruction paths. Removing one witness edge does not guarantee that alternative paths are absent.

Every repair must be verified by rerunning globalization and descent testing.

Step 10: Test transformation stability.

If admissible transformations are supplied, verify for every generating identity edge x Ei y that

g(x) E* g(y).

Because E* is generated by the local edges, checking generators is sufficient to establish global invariance.

For a transformation monoid with generators S, it is enough to check the selected generating transformations if preservation is closed under composition.

The algorithm can also compute quantitative diagnostics.

For every domain Ui, calculate:

the number of local classes,

the number of restricted global classes,

object loss Li,

the number of corrupted global traces,

maximum fusion multiplicity Mi,

and the number of affected states.

For every defect pair or fused class pair, calculate:

shortest witness length,

number of domains traversed,

minimum-confidence edge,

and number of independent witness paths.

For every identity edge, calculate bridge sensitivity by temporarily removing it or by using dynamic connectivity approximations.

A practical system should distinguish three outputs.

Descent certified.

The generated global relation is faithful. The output includes the global classes and a certificate that every local-to-global class map is injective.

Descent obstructed.

The output includes the affected local domains, fused classes, and finite witness paths.

Identity underdetermined.

Descent succeeds, but several global components could be merged without violating current local policies. The minimal globalization is reported together with pairs or components whose identity remains unresolved.

The third output requires analysis beyond basic union-find.

Given minimal global components, one must determine which additional mergers are permitted without creating local collision after transitive closure.

A simple conservative test says that two global components cannot be merged if some domain intersects both in distinct local classes.

Components with no such direct prohibition may still participate in combinations whose joint merger creates an indirect conflict. The full space of faithful globalizations can therefore possess nontrivial combinatorial structure.

For many applications, reporting the minimal globalization is sufficient. It contains all forced identities and no optional ones.

The algorithmic theory also supports streaming updates.

When a new local identity is added, union-find can merge the corresponding components efficiently.

If the merger causes one component to contain two classes from the same domain, a new descent defect appears immediately.

Deletion is harder because removing an edge may split a component. Dynamic graph connectivity methods or periodic recomputation may be required.

This asymmetry mirrors the mathematical closure structure.

Identity addition is monotone.

Identity retraction is reconstructive.

The same applies to conceptual systems.

Adding one equivalence can trigger rapid global fusion.

Undoing that fusion requires determining which identities remain supported by alternative paths.

A provenance-aware implementation is therefore essential.

The system should retain:

which observer supplied each identity edge,

which measurement or rule justified it,

its confidence or admissibility type,

when it was added,

which global mergers depend on it,

and which local distinctions it threatens.

This information transforms the algorithm from a blind clustering procedure into a certifiable discovery system.

The computational examples of this section establish that the local-to-global theory is not only philosophical or categorical. It yields explicit finite tests, efficient algorithms, obstruction certificates, and repair strategies.

The common structure across all applications is now visible.

Local systems form quotient objects.

Overlaps and correspondences connect those objects.

Connected components generate candidate global identity.

Faithful descent prevents a component from containing two locally distinct objects from the same context.

Transformation invariance determines whether the component behaves as one object.

Provenance and admissibility determine whether the resulting fusion should be accepted.

The next section combines these constructions into a distributed object formation theorem and states the resulting research program. It also identifies the limits of exact set-theoretic equivalence and the next required extensions toward graded descent, stacks, policy-selection algebra, and computational discovery.

12. Distributed Object Formation Theorem and Research Program

The preceding sections establish a complete first-order theory of local-to-global object formation. A state space may be observed through multiple restricted domains. Each domain carries an equivalence relation representing its local policy of identity. These policies can agree on every overlap while still failing to arise as the exact restrictions of one global equivalence relation. The obstruction is generated by transitivity: local identities compose across domains and can return to a domain as a new identification that its original policy did not authorize.

The theory separates several operations that are often conflated.

Local observation forms local quotient objects.

Overlap comparison tests direct compatibility.

Equivalence closure constructs the minimal global identity relation forced by all local identifications.

Restriction tests whether that globalization preserves local distinctions.

Transformation invariance determines whether the resulting global classes support representative-independent dynamics.

Categorical descent organizes compatible families, possible global realizations, and the transition from strict quotient objects to higher relational structures.

These operations combine into a distributed object formation theorem. The theorem does not claim that every scientific object is an equivalence class or that observation arbitrarily creates reality. It identifies a general mathematical mechanism by which partial observations can generate effective global objects and specifies exact conditions under which those objects are coherent, faithful, and dynamically stable.

The main conclusion is that distributed objecthood requires more than local agreement.

It requires closure without destructive return.

A local identity chain may extend an object beyond the range of any individual observer. This is the constructive power of globalization. The same chain may reenter a domain and collapse a distinction required there. This is the gluing obstruction.

The theory therefore supplies both a formation law and a failure law.

It explains how local objects become global.

It also explains why some apparently compatible object systems cannot inhabit one common ontology.

The mathematical framework remains intentionally elementary in its foundational layer. Its primitives are sets, subsets, equivalence relations, quotient spaces, complete lattices, graphs, monoid actions, presheaves, and thin categories. Yet these ingredients already produce a nontrivial descent theory. They also identify precisely where richer mathematics becomes necessary.

When strict equality of local policies is too strong, one must pass to equivalence between policies.

When quotienting erases important identity paths, one must retain groupoids.

When local objects glue only up to coherent isomorphism, one must use stacks.

When identity is approximate or uncertain, one must replace exact equivalence with graded, probabilistic, or metric structure.

When the validity of fusion depends on task or transformation, one must make admissibility explicit.

The present section states the combined theorem, summarizes its consequences, defines the emerging research program, and identifies the next mathematical extensions.

12.1 Statement of the Combined Local-to-Global Theorem

Distributed Object Formation Theorem.

Let X be a set and let $\{U_i\}_{i\in I}$ be a cover of X. For each i, let E_i be an equivalence relation on U_i.

Define the direct local identity relation

$R = \bigcup_{i\in I} E_i$.

Define the generated global equivalence

$E^* = \mathrm{EqCl}(R)$.

Then the following statements hold.

First, E* is the unique smallest equivalence relation on X containing every local identity relation Ei.

Second, for x, y ∈ X,

x E* y

if and only if there exists a finite sequence

x = x0, x1, …, xn = y

such that every consecutive pair is related by at least one local policy.

Third, the local family admits a faithful global equivalence relation E satisfying

E|Ui = Ei

for every i if and only if

E*|Ui = Ei

for every i.

Fourth, faithful descent is equivalent to the condition that every generated global equivalence class intersects each domain Ui in at most one Ei-class.

Fifth, faithful descent is equivalent to injectivity of every canonical local-to-global quotient map

ϕi: Ui/Ei → X/E*

defined by

ϕi([x]Ei) = [x]E*.

Sixth, faithful descent is equivalent to the condition that no connected component of the local-object correspondence graph contains two distinct local objects belonging to the same domain.

Seventh, when faithful descent holds, E* is the unique smallest faithful globalization. Every other faithful globalization contains E* and differs from it only through additional identities not determined by the local data.

Eighth, the family of faithful globalizations, when nonempty, is closed under arbitrary intersections and has E* as its least element. It need not be closed under joins and need not possess a greatest element.

Ninth, for every domain Ui, the gluing defect is

$D_i = E^*|_{U_i} \setminus E_i$.

Faithful descent holds exactly when every D_i is empty.

Tenth, every defect pair possesses a finite zig-zag witness connecting two locally distinguished states through locally licensed identities.

Eleventh, let $g: X \to X$ be a transformation. If every local generating relation is mapped into E^*,

$x \; E_i \; y$ implies $g(x) \; E^* \; g(y)$,

then g preserves E^* and induces a unique transformation

$\bar{g}: X/E^* \to X/E^*$.

Twelfth, if a monoid of local transformations agrees on overlaps, assembles into a global action, and preserves every local policy, then it induces a monoid action on the global quotient X/E^*.

Thirteenth, the assignment

$U \mapsto Eq(U)$

with restriction along inclusions forms a contravariant presheaf of complete meet-lattices or thin resolution categories.

Fourteenth, the matching object of a cover consists of all compatible local policy families.

Fifteenth, the restriction map

$\rho: Eq(X) \to Match(\{U_i\})$

need not be injective or surjective.

Failure of surjectivity is gluing obstruction.

Failure of injectivity is global identity underdetermination.

Sixteenth, for a fixed compatible local family e, the fiber $\rho^{-1}(e)$ is its realization category.

The fiber is empty when descent is obstructed.

It contains one element when globalization is unique.

It contains several elements when multiple global ontologies share the same local traces.

Seventeenth, when the realization category is nonempty, E^* is its initial object.

The theorem can be summarized in one sequence:

local equivalence,

union of identity declarations,

equivalence closure,

global quotient formation,

restriction test,

descent certification,

transformation test,

stable distributed objecthood.

The theorem is exact for the stated set-theoretic model.

It does not assume that the local policies are correct, optimal, or scientifically authoritative. It determines what follows from them under equivalence closure and whether their positive identity declarations are compatible with their locally preserved distinctions.

The theorem therefore distinguishes logical consequence from empirical validity.

E* is the identity structure logically forced by the local equivalences.

Whether E* describes the intended objects remains a separate question of admissibility and stability.

12.2 Compatible Local Policies and Minimal Globalization

A compatible family of local policies represents the strongest form of direct observational agreement available before globalization.

For every pair of domains Ui and Uj,

$E_i|_{U_i \cap U_j} = E_j|_{U_i \cap U_j}$.

No shared state pair receives conflicting identity judgments.

This condition is necessary for faithful descent.

It is not sufficient.

The failure of sufficiency is the principal mathematical discovery of the local-to-global theory.

Ordinary intuition often treats overlap agreement as the complete gluing condition. That intuition is correct for many pointwise structures. Compatible local functions, for example, assemble uniquely into a global function. Equivalence relations behave differently because their defining laws include transitive closure.

The local policies do not merely assign independent values to individual states.

They create composable substitutions.

If x may replace y in one domain and y may replace z in another, every global equivalence relation accepting those substitutions must permit x to replace z.

Thus local identity propagates through the observer network.

The minimal globalization

$E^* = \mathrm{EqCl}(\bigcup_i E_i)$

is the exact record of that propagation.

It is canonical.

No order of merging affects it.

No representative choices affect it.

No arbitrary global decisions enter it.

It contains all and only the equivalences forced by finite composition of local identity steps.

The word minimal must be understood relationally.

E^* is the smallest equivalence relation under inclusion.

It is therefore the finest global conceptual resolution accepting every local identity.

It preserves the greatest amount of global distinction compatible with those declarations.

This does not imply that its equivalence classes are small.

A single bridge can join two large components.

Minimality of the relation can produce maximal-looking fusion when the generating network is highly connected.

The global quotient

X/E^*

should therefore be interpreted as the finest unavoidable object space generated by the local observations.

Every map that treats all locally equivalent states identically must factor through X/E^*.

Every other global quotient accepting the local identities is a further quotient of X/E^*.

This gives minimal globalization a universal role.

It is not merely one candidate among many.

It is the common foundation of every globally coherent interpretation that accepts all local positive identity judgments.

The distinction between positive identity and preserved distinction is crucial.

An equivalence declaration is represented directly by inclusion of a pair in Ei.

A distinction is represented by exclusion of a pair from Ei.

Equivalence closure is designed to preserve all included pairs.

It does not automatically preserve excluded pairs.

This asymmetry explains why globalization always exists as an extending relation but may fail as a faithful relation.

The local-to-global problem is therefore a constraint problem with two kinds of information.

Positive constraints require selected states to belong to one global object.

Negative constraints require selected states occurring in one local domain to remain in different objects.

The minimal globalization solves the positive constraints.

The descent test determines whether the negative constraints survive.

The mathematical structure is closely related to consistency in equality theories.

Positive identity declarations generate congruence-like closure.

Local distinctions act as disequality constraints.

A gluing obstruction occurs when the equality closure proves two states equal that one local policy requires to remain unequal.

This reinterpretation supports efficient computation and suggests extensions to automated reasoning.

A discovery system could propose identity edges, compute their closure, and immediately test whether any certified distinction is violated.

Thus minimal globalization is both an algebraic construction and a computational engine for distributed concept formation.

12.3 Faithful Descent and Its Obstruction

Faithful descent requires exact local recovery:

$E^*|U_i = E_i$

for every i.

This condition has a simple geometric meaning.

Take any global object $C \in X/E^*$.

Observe it from a local domain Ui.

Its local trace

$C \cap U_i$

must be either empty or exactly one Ei-class.

A global object may extend beyond Ui.

It may pass through many observers.

It may have no complete local presentation.

It may not return to Ui through a second locally distinct object.

The forbidden return is the gluing obstruction.

This produces a useful image of distributed identity.

An object is a path-connected component in the network of local sameness.

Faithful descent says that the path may travel outward through the network but may not close around a locally preserved separation.

The defect relation

$D_i = E^*|U_i \setminus E_i$

records every such closure.

A pair in Di is not merely an arbitrary disagreement between local and global descriptions. It is a relation forced by all global equivalences accepting the local identity declarations.

The defect is universal.

No coarser global relation can remove it.

No alternative order of gluing can avoid it.

No different choice of quotient representatives can repair it.

The generating system must change.

Several interpretations are possible.

A local identity edge may be erroneous.

A local distinction may be too fine.

Two identity policies may express different tasks and should not be composed.

A shared state may require contextual duplication.

A cover may omit an observer capable of deciding the disputed relation directly.

The equivalence model may be too coarse and should be replaced by a graph, category, or groupoid.

The obstruction itself does not choose among these interpretations.

It identifies the exact location at which one common ontology becomes impossible.

This is already a substantial advance over informal integration.

Instead of saying that two classifications “do not quite fit,” the theory produces:

the affected domain,

the fused local classes,

the global component containing them,

and a finite chain of identity judgments forcing the fusion.

A gluing failure therefore possesses a proof certificate.

The theory also distinguishes obstruction from underdetermination.

Obstruction means no faithful globalization exists.

Underdetermination means more than one faithful globalization exists.

The first arises because local positive identities force a forbidden merger.

The second arises because some cross-domain identities are neither forced nor prohibited.

These are opposite failures of global determination.

An obstructed system contains too much mutually incompatible identity information.

An underdetermined system contains too little information to select one global ontology.

A complete observer network lies between these extremes.

Its local data admit a faithful globalization and determine it uniquely.

The theory thus provides three evaluative questions.

Does a global ontology exist?

If it exists, is it unique?

If it is not unique, which identities are forced and which remain optional?

The minimal globalization answers the third question partially by identifying all forced identities.

The realization category organizes the remaining possibilities.

12.4 Stable Transformation on Global Quotients

Static identity is not sufficient for objecthood.

A mathematical or scientific object must usually persist under a family of relevant transformations. It must be possible to apply those transformations to the object without ambiguity arising from the choice of state representative.

For the generated global equivalence E^*, a transformation $g: X \to X$ descends to X/E^* exactly when

$x \, E^* \, y$ implies $g(x) \, E^* \, g(y)$.

Because E^* is generated by the local relations, this condition can be checked on generators.

It is sufficient that

$x \, E_i \, y$ implies $g(x) \, E^* \, g(y)$

for every i.

Every local identity edge then maps to a global identity path. A finite zig-zag maps to another finite zig-zag. The endpoints remain globally equivalent.

The induced transformation is

$\bar{g}([x]E^*) = [g(x)]E^*$.

This map is representative-independent.

The global quotient can therefore support dynamics, computation, symmetry, intervention, or other operations at the object level.

Distributed transformation adds an assembly problem.

Each observer may possess a local transformation

$g_i: U_i \to U_i$.

For these maps to form one strict global transformation, they must agree on overlaps.

If

$g_i|_{U_i \cap U_j} = g_j|_{U_i \cap U_j}$

for every i and j, then they assemble uniquely into a map $g: X \to X$.

If each g_i preserves E_i, the assembled map preserves E^*.

The result is a stable global quotient.

This theorem is stronger than local stability but weaker than full scientific certification.

A quotient may support well-defined dynamics even when its globalization is not faithful to all local policies.

The dynamics simply respect the coarser generated equivalence.

Conversely, a faithfully descended quotient may fail under a transformation not preserving its classes.

Thus four conditions should be evaluated independently:

local policy validity,

faithful global descent,

functional assembly of transformations,

and invariance of global equivalence.

Their conjunction produces a faithfully stable distributed quotient.

This is the strongest form of objecthood available in the current theory.

Every local object embeds into a global object without local fusion.

Every local transformation participates in one global action.

Every global class has a unique transformed image.

The global object system can therefore support autonomous reasoning.

Transformation invariance also helps distinguish emergent identity from false gluing.

Suppose globalization merges two local objects.

If their states transform coherently under all relevant actions, the fusion may represent a legitimate higher-level object.

If different representatives produce different quotient outcomes, the fusion is destructive.

Stability is therefore a certification test.

It does not prove that the quotient is uniquely correct, but it demonstrates that the proposed identity behaves lawfully under the designated regime.

A future theory should allow the policies themselves to evolve.

In the present framework, E* is fixed while states transform.

In conceptual development, biological emergence, adaptive AI, and theory change, the equivalence relations may refine, coarsen, or reorganize over time.

One then studies trajectories

$E0 \to E1 \to E2 \to \cdots$

in the lattice or category of global resolutions.

Object birth corresponds to refinement.

Object fusion corresponds to coarsening.

Mixed reorganization corresponds to incomparable changes.

Stable distributed objecthood should eventually include both state dynamics on a quotient and policy dynamics among quotients.

The current theorem supplies the fixed-policy foundation required for that extension.

12.5 Presheaf Structure and Sheaf Failure

The assignment

$U \mapsto Eq(U)$

organizes all local observation policies into a presheaf.

An inclusion

$V \subseteq U$

induces a restriction map

$Eq(U) \to Eq(V)$.

Restriction preserves identity and composition.

It preserves arbitrary meets and refinement order.

It does not generally preserve joins because globally generated identity paths may leave V and return.

This failure of join preservation is a lattice-level manifestation of gluing obstruction.

The presheaf viewpoint separates local compatibility from global realizability.

Given a cover {Ui}, the matching object Match({Ui}) contains all families of local policies agreeing on overlaps.

The global restriction map

$\rho: Eq(X) \to Match(\{Ui\})$

sends a global policy to its local traces.

For an ordinary sheaf, ρ would be a bijection.

For conceptual resolutions, it need not be.

The failure of surjectivity means that some compatible local policy families have no global realization.

This is faithful-descent obstruction.

The failure of injectivity means that distinct global identity policies can have exactly the same local restrictions.

This is observational underdetermination.

The sheaf condition therefore fails for two conceptually opposite reasons.

Local data may be globally inconsistent.

Or local data may be globally incomplete.

This conclusion gives the theory a precise relationship to sheaf mathematics without falsely claiming that the equivalence-relation presheaf is already a sheaf.

It also explains why simple sheafification is potentially destructive.

A canonical closure can force a compatible family into global realizability by replacing each Ei with

E*|Ui.

The revised family is globally induced.

But it may be coarser than the original family.

Sheafification in this context can therefore mean ontological revision.

It makes local object systems globally coherent by erasing distinctions that obstruct coherence.

This may be appropriate when the global identity is scientifically justified.

It may be unacceptable when the local distinction is authoritative.

Thus sheafification is not merely formal completion. It embodies a policy choice about which information has priority.

Several refined alternatives arise.

A separation-preserving completion could reject selected identity edges rather than coarsening local policies.

A contextual completion could duplicate states by domain so that identity remains indexed.

A groupoid completion could retain identity paths instead of collapsing them into one class.

A stack-like completion could glue local objects up to coherent equivalence.

A probabilistic completion could preserve uncertainty rather than converting every accepted link into absolute identity.

The ordinary sheaf question is therefore only the beginning.

The deeper issue is to determine the correct kind of global object for the observational data.

Sometimes it is a set quotient.

Sometimes it is a family of quotients.

Sometimes it is a groupoid of presentations.

Sometimes it is a stack of local objects.

The failure of strict sheaf descent is the diagnostic that tells us when the first answer is insufficient.

12.6 Implications for Artificial Intelligence and Science

The theory has direct implications for artificial intelligence because contemporary AI systems routinely construct identities from distributed, heterogeneous, and partial information.

A multimodal system may integrate images, text, audio, sensor data, and structured records.

Each modality defines its own local distinctions.

Cross-modal alignment supplies identity links.

Equivalence closure can then generate global concepts or entities.

The theorem shows why pairwise alignment is not enough.

Every individual correspondence may be plausible while the resulting global component contains concepts that one modality must distinguish.

A multimodal AI system should therefore test descent at the component level.

It should not merely optimize local matching scores.

It should verify that no global concept cluster contains incompatible local classes.

This requirement applies to entity resolution, retrieval, memory consolidation, knowledge graphs, model merging, federated learning, agent communication, and multimodal representation.

AI memory provides a particularly important application.

A system may store several observations as memories of one entity. Local cues link new memories to old ones. Transitive closure can gradually expand the identity cluster.

One mistaken link may contaminate an entire memory object.

A descent-aware memory system would retain:

the local contexts in which identity was asserted,

the paths connecting memories,

the distinctions threatened by a merger,

and the transformations under which the merged memory remains coherent.

It would also distinguish forced identity from optional identity.

This could prevent the premature consolidation of ambiguous experiences.

Federated and multiagent AI systems face the same problem at larger scale.

Each agent develops local categories.

Communication creates correspondences.

A global ontology emerges only if the correspondences do not create destructive returning paths.

The theory supplies an exact audit.

For science, the implications are equally broad.

Scientific objects are often assembled across instruments, laboratories, databases, scales, and theoretical languages.

A disease entity may combine symptoms, biomarkers, imaging patterns, genomic variants, and treatment responses.

An astronomical object may combine observations across wavelengths and epochs.

A biological lineage may combine samples collected across regions.

A climate structure may combine local measurements from many sensors.

A particle or field object may be inferred from different experimental channels.

A species, ecosystem, storm, or geological formation may never be observed completely by one instrument.

These are distributed objects.

Their unity is generated through networks of local identity.

The theory imposes a discipline on such construction.

The local policies must be stated explicitly.

Their overlaps must be tested.

The global equivalence closure must be computed or approximated.

Returning fusions must be identified.

The resulting classes must be tested under relevant transformations.

This procedure can improve scientific reproducibility.

Disagreements often arise because different groups use different implicit identity policies. One laboratory treats two samples as equivalent under one criterion while another preserves a distinction relevant to a different task.

Making the policies explicit reveals whether the dispute concerns data, observation, transformation, or ontology.

The theory also reframes interdisciplinary integration.

Two disciplines may use locally successful concepts connected by partial translation. A chain of translations can produce a false identification at the original endpoint. The problem is not necessarily that any translation is wrong. The composition may fail to preserve the relevant identity policy.

This is conceptual drift represented mathematically.

A descent-aware integration would track which equivalences may compose and which distinctions must remain protected.

The theory therefore supports a broader conception of scientific objectivity.

Objectivity need not require one observer with complete access.

It can arise from distributed observations when the local identity system descends faithfully and the global quotient remains stable under admissible transformations.

Objectivity becomes certified coherence across observers.

This does not make objecthood arbitrary or merely consensual.

The network is constrained by equivalence laws, overlap data, negative distinctions, dynamics, and empirical tests.

A global object is accepted because it survives these constraints.

The theory also gives mathematical form to emergence.

An emergent object can be a global equivalence class that no local observer sees completely but that is assembled through faithful or scientifically justified gluing.

Its reality lies in stable distributed coherence.

This applies naturally to organisms, institutions, ecosystems, languages, markets, and other systems whose unity is not localized in one observation.

The same formalism also warns against overintegration.

Not every connected data cluster is one object.

Not every chain of semantic similarity supports identity.

Not every locally valid abstraction can be composed globally.

Object formation requires admissibility, not connectivity alone.

12.7 Limits of Set-Theoretic Local Identity

The exact theory developed in this paper is powerful because it is simple. That simplicity also defines its limits.

The first limitation is binary identity.

An equivalence relation declares two states either equivalent or not equivalent.

Many observations support degrees of similarity, confidence, or compatibility rather than exact sameness.

A probabilistic match of 0.8 cannot be represented faithfully as an unconditional equivalence edge without losing uncertainty.

The second limitation is unrestricted transitivity.

Equivalence requires that every identity chain compose completely.

Similarity, influence, translation, and biological relatedness often weaken over paths.

A sequence of strong local similarities may connect globally dissimilar endpoints.

Treating similarity as equivalence creates false gluing.

The third limitation is loss of provenance.

The quotient X/E* records which states belong to one class but not why they were identified.

It forgets which observer supplied each relation, which measurement supported it, and which paths connected the states.

The fourth limitation is loss of path multiplicity.

Several independent identity paths may provide stronger evidence than one path. The quotient treats them identically.

Different paths may also support different interpretations. This information disappears after component collapse.

The fifth limitation is contextual ambiguity.

A state may play different roles in different domains.

Representing it as one literal element of X automatically joins those contexts.

A more faithful model may require domain-indexed copies and explicit comparison morphisms.

The sixth limitation is exact overlap equality.

Two local policies may be structurally equivalent without using the same labels, state representations, or literal relations.

Strict descent rejects such systems even when an isomorphism connects them.

The seventh limitation is fixed ontology.

The theory treats the local and global equivalence relations as supplied policies.

It does not yet model how a system chooses them, learns them, or revises them in response to evidence.

The eighth limitation is deterministic transformation.

The induced quotient maps assume functions. Many systems involve nondeterministic, stochastic, branching, or relational dynamics.

The ninth limitation is set-level objecthood.

Some objects are better represented by categories, groupoids, complexes, manifolds, measures, or processes than by equivalence classes.

The tenth limitation is absence of higher coherence.

When local objects glue through isomorphisms, one must track compatibility among those isomorphisms on triple overlaps and transformations among compatibility witnesses.

These limitations define the route forward.

Exact equivalence should be retained as the foundational case because it supplies clear theorems, efficient algorithms, and interpretable obstruction certificates.

The next structures should extend rather than discard it.

A weighted identity graph can attach confidence to local links.

A fuzzy equivalence can represent graded sameness.

A metric or Lawvere-enriched category can represent cost or distortion.

A probabilistic kernel can represent uncertain transformation.

A groupoid can retain states and reversible identity witnesses.

A category can retain noninvertible abstraction and translation.

A stack can support local objects gluing up to coherent isomorphism.

A higher stack can retain transformations among gluing equivalences.

A policy-indexed fibration can separate identity regimes by task, scale, or observer.

The exact quotient X/E* then appears as one decategorified or thresholded output of a richer system.

The central methodological principle should be:

Do not collapse more structure than the application can safely forget.

When paths, contexts, confidence, or provenance matter, retain them.

Quotient only after admissibility has been established.

This principle connects the present theory directly to the broader algebra of object formation.

Stable quotienting is not merely a way to simplify a state space.

It is a certification that selected differences may be forgotten.

Distributed descent adds a second certification.

It asks whether those acts of forgetting remain mutually coherent across observers.

A mature theory of objecthood must satisfy both.

12.8 Next Steps: Graded Descent, Stacks, and Computational Discovery

The next stage of the research program should proceed through several theorem-bearing extensions.

The first is graded descent.

Replace exact local equivalence relations with weighted or graded identity relations.

A local observer may assign a degree

w_i(x, y)

representing confidence, similarity, substitution cost, or evidential strength.

The global problem is no longer ordinary connectedness. A long path of moderate links should not necessarily create absolute identity.

Possible globalization rules include:

maximum bottleneck strength,

minimum path cost,

probabilistic combination,

bounded-length closure,

and semiring-valued path composition.

The main mathematical question becomes:

Under what conditions do locally graded identity relations produce a globally coherent graded object structure without amplifying weak evidence into false certainty?

This should lead to a graded faithful descent theorem.

The second extension is typed identity.

Different equivalence edges may represent different forms of sameness:

visual identity,

causal identity,

functional identity,

legal identity,

semantic equivalence,

or shared origin.

Not every type should compose with every other type.

A typed identity algebra would specify admissible compositions.

The resulting closure would be generated not by unrestricted transitivity but by a grammar or category of licensed paths.

False gluing would then appear as an inadmissible path type rather than only as a violated local distinction.

The third extension is groupoid globalization.

Instead of collapsing every connected component into one quotient point, retain states as objects and local equivalences as invertible morphisms.

The path groupoid records identity witnesses and their composition.

The set quotient X/E* is recovered as the set of connected components.

This extension preserves provenance and multiple paths.

A central theorem should characterize when the groupoid can be safely decategorified without loss of task-relevant structure.

The fourth extension is stack-theoretic observation.

Local object categories should be assigned to domains.

Overlap agreement should occur through equivalences of categories.

Triple overlaps should carry coherence isomorphisms.

The main problem becomes effective descent of local objects and morphisms.

This will allow the theory to represent objects having multiple local presentations that are equivalent but not literally identical.

The fifth extension is policy-indexed descent.

Identity should be indexed by an admissibility policy π.

For each policy, one has a relation Eπ.

Policies may form a lattice, category, or bicategory.

Globalization should then preserve not one universal notion of sameness but a structured family of identities.

A pair may descend under one policy and remain obstructed under another.

The central question becomes:

How do descent, refinement, and transformation vary functorially with policy?

This would connect local-to-global object formation with Admissibility Algebra.

The sixth extension is dynamic descent.

Let the domains, policies, and identity links vary with time.

One obtains a sequence or flow of observer networks

Nt.

The generated global relations E*t can split, merge, appear, and disappear.

Object persistence becomes correspondence among quotient classes across changing policies.

Formation thresholds, extinction, bifurcation, and phase transitions can then be studied as dynamics on the space of global resolutions.

The seventh extension is descent under stochastic transformation.

Local transitions may be Markov kernels rather than functions.

A quotient is stable when equivalent states induce the same or appropriately equivalent probability distributions on quotient objects.

This connects distributed object formation with probabilistic bisimulation, lumpability, and stochastic abstraction.

The eighth extension is automated obstruction repair.

Given weighted identity declarations and authoritative separation constraints, compute a maximally informative global object system satisfying descent.

Possible objectives include:

maximizing accepted identity weight,

minimizing deleted links,

minimizing local object loss,

preserving selected high-priority domains,

or optimizing transformation stability.

This becomes a constrained optimization problem over equivalence relations or richer identity structures.

The ninth extension is active observer design.

When global identity is obstructed or underdetermined, determine which new observation would be most informative.

A new domain may contain selected states from different components.

Its policy can decide whether they should be fused or separated.

The design problem asks:

Which additional overlap or measurement resolves the largest amount of uncertainty at the lowest cost?

This turns descent theory into an experimental design calculus.

The tenth extension is computational concept discovery.

A discovery system could begin with a state space, candidate observations, transformations, and protected distinctions.

It could propose local equivalence relations, compute their global closure, detect defects, and revise the policies.

The loop would be:

generate candidate local identities,

construct minimal globalization,

test faithful descent,

test transformation invariance,

measure compression and predictive value,

repair or revise the policies,

and retain certified quotient objects.

This is not ordinary clustering.

Clustering groups data according to a selected similarity criterion.

Computational object discovery must determine which similarities support stable, composable, and policy-admissible identity.

The theory developed in this paper provides the exact logical core of such a system.

The final research program can therefore be expressed as a progression.

Exact local equivalence establishes the foundational theorem.

Gluing obstruction identifies the failure of naive local compatibility.

Stable globalization adds lawful dynamics.

Presheaf analysis organizes restriction and underdetermination.

Groupoids retain identity witnesses.

Stacks support gluing up to equivalence.

Graded structures represent uncertainty and distortion.

Policy indexing controls admissible substitution.

Dynamic descent models changing ontologies.

Computational discovery searches the resulting space of possible object systems.

The central mathematical problem remains constant throughout these extensions:

When may many partial presentations be treated as one object without destroying a distinction required somewhere in the system?

The present paper answers that question for exact equivalence relations.

The answer is not simply overlap agreement.

It is faithful closure.

Local identities form a coherent global object system exactly when every identity generated by their transitive composition is already licensed in every domain capable of comparing its endpoints.

This is the local-to-global law of object formation.

It extends stable quotienting from one observer to networks of observers.

It explains both the emergence of distributed objects and the failure of false fusion.

It gives object formation a descent theory.

And it establishes a rigorous path from the algebra of conceptual resolution toward sheaves, stacks, higher identity, and computational discovery.

13. References

1. Youvan, Douglas C. Object Formation and Conceptual Resolution: Stable Quotients, Complete Refinement Lattices, and Functorial Observation. July 2026. DOI: 10.13140/RG.2.2.25921.72800.
2. Youvan, Douglas C. The Mathematics of Concept Formation: Discovery Calculus, Explanation Geometry, Theory Dynamics, and the Formal Birth of New Objects. July 2026. DOI: 10.13140/RG.2.2.23562.43209.
3. Youvan, Douglas C. How Mathematics Makes Objects: Observation, Identity, and a New Algebra of Concept Formation. July 2026. DOI: 10.13140/RG.2.2.13338.81603.
4. Youvan, Douglas C. Mathematical It from Bit: Stable Quotients and the Algebra Before Abstract Algebra. July 2026.
5. Abramsky, Samson, and Adam Brandenburger. “The Sheaf-Theoretic Structure of Non-Locality and Contextuality.” New Journal of Physics, vol. 13, 2011, article 113036.
6. Anderson, P. W. “More Is Different: Broken Symmetry and the Nature of the Hierarchical Structure of Science.” Science, vol. 177, no. 4047, 1972, pp. 393–396.
7. Ashby, W. Ross. An Introduction to Cybernetics. Chapman and Hall, 1956.
8. Awodey, Steve. Category Theory. 2nd ed. Oxford University Press, 2010.
9. Baier, Christel, and Joost-Pieter Katoen. Principles of Model Checking. MIT Press, 2008.
10. Bar-Shalom, Yaakov, X. Rong Li, and Thiagalingam Kirubarajan. Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software. Wiley, 2001.
11. Bénabou, Jean, and Jacques Roubaud. “Monades et descente.” Comptes Rendus de l’Académie des Sciences de Paris, Série A, vol. 270, 1970, pp. 96–98.
12. Bhattacharya, Indrajit, and Lise Getoor. “Collective Entity Resolution in Relational Data.” ACM Transactions on Knowledge Discovery from Data, vol. 1, no. 1, 2007, article 5.
13. Birkhoff, Garrett. Lattice Theory. 3rd ed. American Mathematical Society, 1967.
14. Blackman, Samuel S., and Robert Popoli. Design and Analysis of Modern Tracking Systems. Artech House, 1999.
15. Blackwell, David. “Equivalent Comparisons of Experiments.” Annals of Mathematical Statistics, vol. 24, no. 2, 1953, pp. 265–272.

16. Borceux, Francis. Handbook of Categorical Algebra. Vol. 1: Basic Category Theory. Cambridge University Press, 1994.

17. Borceux, Francis. Handbook of Categorical Algebra. Vol. 3: Categories of Sheaves. Cambridge University Press, 1994.

18. Bredon, Glen E. Sheaf Theory. 2nd ed. Springer, 1997.

19. Buchholz, Peter. “Exact and Ordinary Lumpability in Finite Markov Chains.” Journal of Applied Probability, vol. 31, no. 1, 1994, pp. 59–75.

20. Christen, Peter. Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection. Springer, 2012.

21. Cousot, Patrick, and Radhia Cousot. “Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints.” Proceedings of the Fourth ACM Symposium on Principles of Programming Languages, ACM, 1977, pp. 238–252.

22. Davey, B. A., and H. A. Priestley. Introduction to Lattices and Order. 2nd ed. Cambridge University Press, 2002.

23. Doan, AnHai, Alon Halevy, and Zachary Ives. Principles of Data Integration. Morgan Kaufmann, 2012.

24. Eilenberg, Samuel, and Saunders Mac Lane. “General Theory of Natural Equivalences.” Transactions of the American Mathematical Society, vol. 58, no. 2, 1945, pp. 231–294.

25. Euzenat, Jérôme, and Pavel Shvaiko. Ontology Matching. 2nd ed. Springer, 2013.

26. Fellegi, Ivan P., and Alan B. Sunter. “A Theory for Record Linkage.” Journal of the American Statistical Association, vol. 64, no. 328, 1969, pp. 1183–1210.

27. Fong, Brendan, and David I. Spivak. An Invitation to Applied Category Theory: Seven Sketches in Compositionality. Cambridge University Press, 2019.

28. Ganter, Bernhard, and Rudolf Wille. Formal Concept Analysis: Mathematical Foundations. Springer, 1999.

29. Gärdenfors, Peter. Conceptual Spaces: The Geometry of Thought. MIT Press, 2000.

30. Getoor, Lise, and Ashwin Machanavajjhala. “Entity Resolution: Theory, Practice and Open Challenges.” Proceedings of the VLDB Endowment, vol. 5, no. 12, 2012, pp. 2018–2019.

31. Giraud, Jean. Cohomologie non abélienne. Springer, 1971.

32. Goldenfeld, Nigel. Lectures on Phase Transitions and the Renormalization Group. Addison-Wesley, 1992.

33. Goodfellow, Ian J., Jonathon Shlens, and Christian Szegedy. “Explaining and Harnessing Adversarial Examples.” International Conference on Learning Representations, 2015.

34. Grätzer, George. General Lattice Theory. 2nd ed. Birkhäuser, 1998.

35. Grothendieck, Alexander. Revêtements étales et groupe fondamental. Séminaire de Géométrie Algébrique du Bois Marie 1960–1961, SGA 1. Lecture Notes in Mathematics, vol. 224, Springer, 1971.

36. Haken, Hermann. Synergetics: An Introduction. 3rd ed. Springer, 1983.

37. Hall, David L., and James Llinas. “An Introduction to Multisensor Data Fusion.” Proceedings of the IEEE, vol. 85, no. 1, 1997, pp. 6–23.

38. Hollander, Sharon. “A Homotopy Theory for Stacks.” Israel Journal of Mathematics, vol. 163, 2008, pp. 93–124.

39. Janelidze, George, and Walter Tholen. “Facets of Descent, I.” Applied Categorical Structures, vol. 2, 1994, pp. 245–281.

40. Johnstone, Peter T. Sketches of an Elephant: A Topos Theory Compendium. Vols. 1–2. Oxford University Press, 2002.

41. Kanellakis, Paris C., and Scott A. Smolka. “CCS Expressions, Finite State Processes, and Three Problems of Equivalence.” Information and Computation, vol. 86, no. 1, 1990, pp. 43–68.

42. Kashiwara, Masaki, and Pierre Schapira. Categories and Sheaves. Springer, 2006.

43. Kemeny, John G., and J. Laurie Snell. Finite Markov Chains. Van Nostrand, 1960.

44. Kevrekidis, Ioannis G., C. William Gear, James M. Hyman, Panagiotis G. Kevrekidis, Olof Runborg, and Constantinos Theodoropoulos. “Equation-Free, Coarse-Grained Multiscale Computation: Enabling Microscopic Simulators to Perform System-Level Analysis.” Communications in Mathematical Sciences, vol. 1, no. 4, 2003, pp. 715–762.

45. Larsen, Kim G., and Arne Skou. “Bisimulation through Probabilistic Testing.” Information and Computation, vol. 94, no. 1, 1991, pp. 1–28.

46. Lenzerini, Maurizio. “Data Integration: A Theoretical Perspective.” Proceedings of the Twenty-First ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, ACM, 2002, pp. 233–246.

47. Mac Lane, Saunders. Categories for the Working Mathematician. 2nd ed. Springer, 1998.

48. Mac Lane, Saunders, and Ieke Moerdijk. Sheaves in Geometry and Logic: A First Introduction to Topos Theory. Springer, 1992.

49. Mahler, Ronald P. S. Statistical Multisource-Multitarget Information Fusion. Artech House, 2007.

50. Medin, Douglas L., and Edward E. Smith. “Concepts and Concept Formation.” Annual Review of Psychology, vol. 35, 1984, pp. 113–138.

51. Milner, Robin. Communication and Concurrency. Prentice Hall, 1989.

52. Newcombe, Howard B., James M. Kennedy, S. J. Axford, and A. P. James. “Automatic Linkage of Vital Records.” Science, vol. 130, no. 3381, 1959, pp. 954–959.

53. Ore, Oystein. “Theory of Equivalence Relations.” Duke Mathematical Journal, vol. 9, no. 3, 1942, pp. 573–627.

54. Paige, Robert, and Robert E. Tarjan. “Three Partition Refinement Algorithms.” SIAM Journal on Computing, vol. 16, no. 6, 1987, pp. 973–989.

55. Pavliotis, Grigorios A., and Andrew M. Stuart. Multiscale Methods: Averaging and Homogenization. Springer, 2008.

56. Pearl, Judea. Causality: Models, Reasoning, and Inference. 2nd ed. Cambridge University Press, 2009.

57. Riehl, Emily. Category Theory in Context. Dover Publications, 2016.

58. Rosch, Eleanor. “Principles of Categorization.” In Cognition and Categorization, edited by Eleanor Rosch and Barbara B. Lloyd, Lawrence Erlbaum Associates, 1978, pp. 27–48.

59. Shannon, Claude E. “A Mathematical Theory of Communication.” Bell System Technical Journal, vol. 27, 1948, pp. 379–423 and 623–656.

60. The Stacks Project Authors. The Stacks Project. Open mathematical reference work.

61. Vistoli, Angelo. “Grothendieck Topologies, Fibered Categories and Descent Theory.” In Fundamental Algebraic Geometry: Grothendieck’s FGA Explained, edited by Barbara Fantechi, Lothar Göttsche, Luc Illusie, Steven L. Kleiman, Nitin Nitsure, and Angelo Vistoli, American Mathematical Society, 2005, pp. 1–104.

62. Wilson, Kenneth G., and J. Kogut. “The Renormalization Group and the ε Expansion.” Physics Reports, vol. 12, no. 2, 1974, pp. 75–199.